SELECTED PAPERS

©1965-77 Theodor H. Nelson

TO THE READER

Ideally, there should be a complete introduction to the papers herein, clarifying them and explaining their contexts of origin. Perhaps there will be a later, published version for which that can be done. For the present all I can say is that these papers represent my thinking on how computer systems should help man's (and woman's) thought. They should be simple, clear and highly interactive, helping us with our most complicated mental tasks.

The hypertext concept is obvious. The thinkertoy concept may not seem obvious at first. Think yourself to a world of only screens, though, and keeping track of what you are looking at and thinking about becomes the fundamental problem.

The collateration concept, discussed more recently in Computer Lib, is my basic proposal of a way to structure our thinkertoys. I am eager to hear any others.

I write this at the brink of a new world. The next years-- the next months-- are going to be like nothing that has ever happened before. The impact of home computers will be like the impact of television-- but far more beneficial, I hope, for our mental life.

This is going to be good.

Ted Nelson

Swarthmore College

February, 1977

VERY GENERAL

Handout (1973)

New Media and Creative Facilities (1967?)

Barnum-Tronics (1970)

THINKERTOYS AND THE COLLATERATION CONCEPT; THE HYPERTEXT NOTES
(The Hypertext Notes actually combine all the topics, but they are together here for simplicity's sake.)

Decision/Creativity Systems (1970)
The basic Thinkertoy concept.

Hypertext Notes 0-10 (1967)
Thinkertoys and hypertexts generally considered; the basic concept of the Xanadu[tm] Information System.

Hypertext Note 102 (1976)
(This is actually a companion piece to the final paper.)

A File Structure for the Complex, the Changing and the Indeterminate (1965)
The basic idea of collateration, interwoven with a number of other things.

HYPERTEXT, INFORMATION RETRIEVAL AND COMPUTER-ASSISTED INSTRUCTION

As We Will Think (1968)
A new look at Vannevar Bush's 1945 "memex" system: what Information Retrieval theorists have overlooked in the fundamental writing.

Getting It Out of Our System (1967)
More on Information Retrieval versus hypertext. This contains the interesting analogy between computer display and movies.

Computopia and Cybercrud (1971)
More on Computer-Assisted Instruction versus hypertext. This continues the ideas of "No More Teachers' Dirty Looks" (in C. Lib).

HIGHLY INTERACTIVE SYSTEMS AND WINDOWING

A Conceptual Framework for Man-Machine Everything (1973)
This basically introduces a terminology for discussing highly interactive computer-screen systems. (A severely altered version appeared in the proceedings of the National Computer Conference for that year.)

Data Realms and Magic Windows (1975)
Extension of basic concepts to business and other interactive systems.

Design of a Transcendental Literary Network (1977)
This (along with Hypertext Note 102) describes the crystallized design proposed for the Xanadu Literary Network.

HANDOUT

on Computer Arts & Media

Theodor H. Nelson
Lecture of 29 Jan 73

PRELIMINARY MATERIAL,
NOT FOR PUBLICATION
but © 1973 T. Nelson
on general principles.

COMPUTER GRAPHICS & HYPERSYSTEMS are one of the Grandest Developments in the History of Civilization. Started around 1955, have not yet hit academic & consumer markets on widespread basis. Just wait.

Conservative view of impact: graphic terminals will revolutionize large organizations & business world. I think, however, real impact will be on PERSONAL & HOME USE, beginning around 1975, and wholly restructuring our use of the written word, etc., by around 1980. Nobody is prepared for this, least of all the computer industry.

THE SCREENS ARE COMING!
--display screens on which we will read, write and play. Some types (overlapping):

CATHODE RAY TUBE (CRT)-- like that in a TV set but responding to you. Capable of animation.

STORAGE TUBE-- cathode-ray tube which holds an image (non-moving) that has been painted on it just once.

PLASMA PANEL-- has many tightly-packed dots which light up or don't. Flat.

"Terminal"-- requires a mother computer feeding it.

Stand-alone system-- complete in itself.

Keyscope-- like a typewriter but with a screen instead of paper, usually text only.

Line graphic console-- shows line-drawings on the screen (possibly animated).

Halftone graphic system-- shows grey-level or areas of color.

3D halftone graphic system-- shows imaginary "things" as if in a photograph.

In other words, that makes possible MOVIE SYNTHESIZERS, as in the systems by
Gen. Electric,
U. of Utah/Evans & Sutherland Computer Co.,
and MAGI (Mathematical Applications Group, Inc.).

My own system, Fantasm™ (not yet build but pat. pending) is supposed to animate Hieronymus Bosch, simulating real photography of incredibly complex scenes. Or is that just a lot of Bosch?

INTERACTIVE GRAPHIC SYSTEMS-- allow you to play with pictures already stored, or draw your own-- in 2D or 3D. "Computer-Aided Design" is one of the catchwords here. Means almost anything.

INTERACTIVE TEXT SYSTEMS speed up work incredibly. Vannevar Bush predicted it all in '45 but everybody pays lip service to the article without reading it. Doug ENGELBART, of Stanford Research, has an incredible setup, hundreds of guys reading and writing on screens. My XANADU™ is the minicomputer equivalent; three Xanadus are a network by themselves. Every home by 19xx.

COMPUTER-ASSISTED INSTRUCTION as an ideal seems to me to be a mistake.
a) at the low, or automated workbook level, it does not foster initiative or mental independence (though it does make drudgery more pleasant and effective);
b) at the high level, the ideal of a godlike Teacher who converses with you-- A) it's too hard to do to be practical, B) who wants systems that talk back?

Philosophical alternative: user control of rich Responding Resources.

BASIC CONCEPTS

COMPUTER-- automatic device which may be preloaded to make it behave intricately.

COMPUTER LANGUAGE-- spell-casting system to preload computer for a specific range of purposes. A good programmer can pick up new systems of spell-casting very quickly.

DATA STRUCTURES-- the real issue in all computer system design: what kinds of info & interrelation are stored how in the computer.

TURNKEY SYSTEM, or "user-level" system, requires no further programming, just behaves in some Canned fashion. Usually for "naive user."

INPUT vs PROGRAMMING-- turnkey systems, e.g. for text & linked ideas, won't have to be programmed; that's already done & hidden. You just input, e.g. type & manipulate screen controls.

FANTIC CONTROLS vs COMPUTER LANGUAGES -- similarly, use of the hot systems will be through controls with a sportscar feel. (See below.)

BIG COMPUTER VIEW

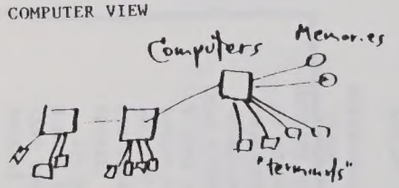

Future info. service system will be lots of Big Computers hooked together by high-speed lines, plus lots of little terminals attached to the big computers like piglets.

MINICOMPUTER NETWORK VIEW

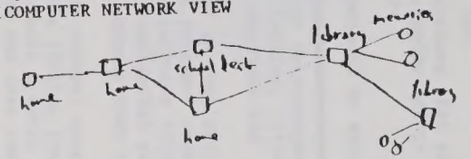

We'll all have our own minis with screens and memories in a network; phooey on big machines.

HYDROT

OVERALL CONCEPTS

FANTICS-- Art & Science of Presentation. Thus Writing, Diagramming & Movie-making are branches of Fantics. Key conjecture: computers will be involved in all branches of fantics by 1980. LOYAL TO BERT. RUSSELL'S DICTUM: "The reason is, and by rights ought to be, slave to the emotions." Pretending it's cold technology is always a fake.

FANTIC SPACE
The "space" or spatial conceptual structure of a fantic system or presentation. (Cf. "filmic space.")

FANTIC FIELD SYSTEM
System whose operation generates a complete conceptual/manipulative environment.

FANT-- general-purpose interactive system for text and pictures. E.g., PLATO? XANADU?

FANTIC CONTROLS
Controls having conceptual unifying aspect-- e.g. joysticks but not gauges.

Hyper-Medias.

HYPERTEXTS: non-sequential writing. E.g.

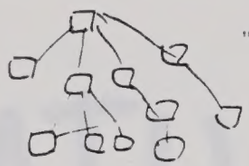

Discrete "chunk style" H'text with various branches available;

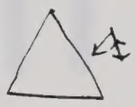

"Stretchtext"™ which gets longer or shorter

Quoteback facilities: if you want to see where a quote came from, press a button and its surroundings are replaced by the original; and so on.

HYPERGRAPHICS is then the creation of multidimensional explorable pictures, those that respond, etc.

Hypercomics: you may choose what the characters do next. Etc.

HYPERSYSTEMS are then systems that allow the user to explore complex information in his own way.

HYPER-MEDIA: Multidimensional or extended (hyper-) media.

MOOTING SYSTEMS allow the intercomparison of complex theories, designs, cities, versions of a plan, uncertain and undecided things.

Thus a CREATIVITY SYSTEM is some sort of a text and/or picture system-- preferably both-- with mooting structures allowed, so you can try out very complicated plans and arrangements and compare them in detail.

(It may also employ control systems, e.g. a high-power movie editing system that lets you rapidly edit a movie several different ways and then compare.)

Zippered Collateral Structures are my notion of a way to handle mooting systems, text annotations and file couplings in general. This is a generalization of side-by-side printing, allowing many parts of objects to be coupled one-for-one any-which-way among complex objects. A "multicoupler" holds the parts together.

Deep Stuff.

Thus we can mix-n-match the interactive COMPUTER ARTS & MEDIA:

(HYPER-) (nonsequential) ± Text* ± 2D 3D etc } GRAPHICS (line or half-tone) / motion, maybe } ± ANIMATION, visual overlays, etc ± MULTI- COUPLING ± FANTIC (physical) controls ± 3-D viewing simulation, textures, etc.

I do not claim this to be exhaustive, but I do think these are the central concepts. Computer people tend to find this shallow: these data structures are comparatively simple, etc.

But the point is: WHAT DOES THE HUMAN MIND WANT? For-- scholarship, learning, creativity, fun } which are intrinsically entwined, Thanks.

⇨ WHAT ARE THE CLEVEREST AND BEST UNIFICATIONS?
How resolve simplicity, power and conceptual unity?
E.g. Parallel Textface™, Walking Net™.

Summary: WHY COMPUTER MEDIA?

1. Non-sequential organization possible: HYPER-MEDIA
   Readers/Users/Viewers may now browse among nonsequentially organized Anything
2. Text & graphics-- 2D or 3D-- easily stored, revised, transmitted.
3. Mooting systems, annotative facilities lend new clarity to thoughtful work.

WHITHER CIVILIZATION ANYHOW?

It follows that making this world (of computers and thought) what we would <u>like</u> it to be-- especially wrenching it from those who claim it "has to be" a certain way-- is one of the more pressing problems of this historical instant.

VICTORIAN OUTLOOK OF PRESENT AUTHOR

I think that individual authorship and the Grand Corpus of our written heritage are precious values; and that the new age of hypertext should build on these traditions, rather than mush us into committee authorship and indifference to the past.

BUT
let us forge directly toward the Screen Future, and the creation of screenworlds we will love to live in.

ALSO
remembering the problems of PRIVACY, ACCESS for everybody, What Gets Kept?, DANGER to the corpus-- the more we store digitally, the safer it is from book-damage but the more liable to certain other problems.

Let us hasten thoughtfully.

Digitized by the Internet Archive
in 2025

https://archive.org/details/selectedpapers1900tedn

(5)

ARCHIVE
ARCHIVE

NEW MEDIA AND CREATIVE FACILITIES

NEW MEDIA
AND
CREATIVE FACILITIES

NEW MEDIA

ADVANCED PRESENTATIONAL SYSTEMS: Public or publishable hyper-media.

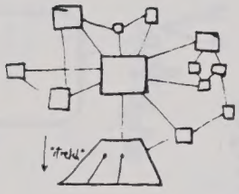

HYPERTEXTS

New forms of writing, to be read from a computer display screen, for learning or search.

With sources, details and expository consecutively available, the reader may select and pursue with unprecedented freedom, choosing what part of the whole to see next, and how it shall look and behave.

Hypergrams

Visual presentations on a computer display screen which change according to the viewer's preference and manipulations

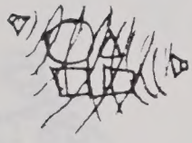

MEDIA MIX

multi-screen multi-amplifier and branching exhibits pavilions environments and simulations

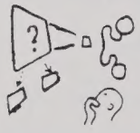

HYPERFILMS AND BRANCHING MOVIES

These problems are coordinated and unified; the two types of system are the same. The difference is only

These systems are, with varying emphases, complex presentational environments embracing both

Design of their presentational matter and manipulation requires

THE COMMON PROBLEMS. Structuring graphics of information, text or

movie sequences, with their complex overlays and links.

CREATIVE FACILITIES

MACHINES TO AID THE MIND: Studios for private generating, considering, organizing.

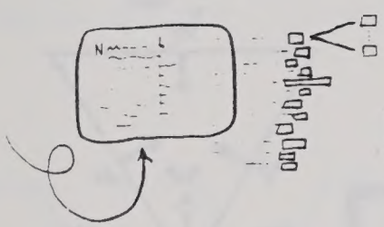

TEXT CONSOLES
for input, proofing, update and query:
HIGH-POWER TEXT CONSOLES
where lurk complex interlinked text structures of user design, on call, indefinitely revisible any which way.
AIDS TO THOUGHT:
systems for the annotation and dissection of problems and ideas, consequences, discords and futures.

FILM & VIDEO EDITING CONSOLES
Automatic switching and transfer with annotation and aids to considering.

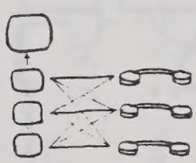

MOVIE PRODUCTION CONSOLES
control of: camera movement, video playback, script and production notes.

INTRINSIC PRODUCTION CONSOLES
With complete script handling, annotation of data, and filing of partial results

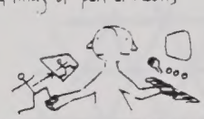

FOR ANIMATION:
Titles, story boards, motor and full production — low-cost or high-quality

FOR AUDIO MANIPULATION:
editing, modulation, effects management and synthesis: with simplified controls, and direct command of effects

PICTURE-EDITING, INDEXING AND LAYOUT CONSOLES
for magazines, exhibitions and archives — gathering materials for comparative effects, annotating and filing

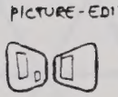

...hen such systems are used for presentation, the material need be revised less often.

...some function, organizing both and making them vivid and tangible.

...e of human imagination, impressions, thought and decision.

...ic interconnections of materials and data (each volume I&O;

...ter as dashboard environment crafting.

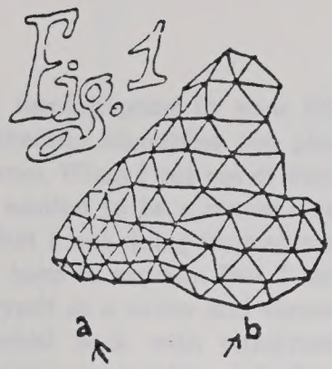

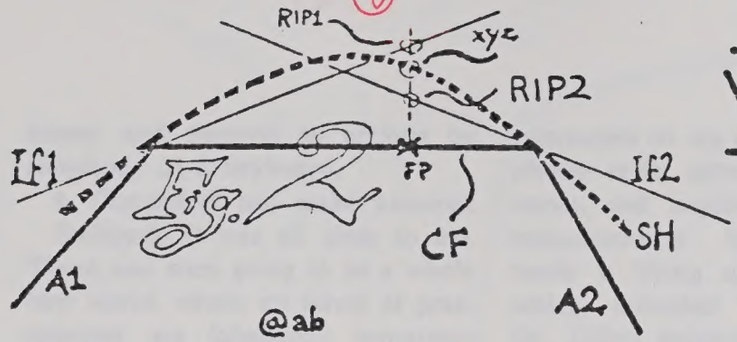

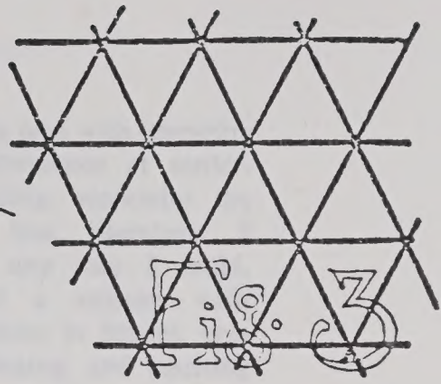

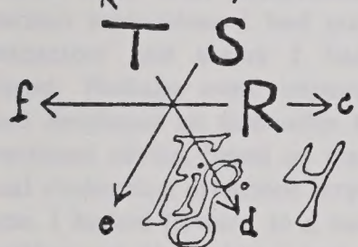

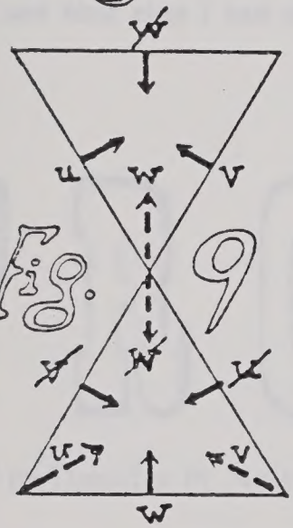

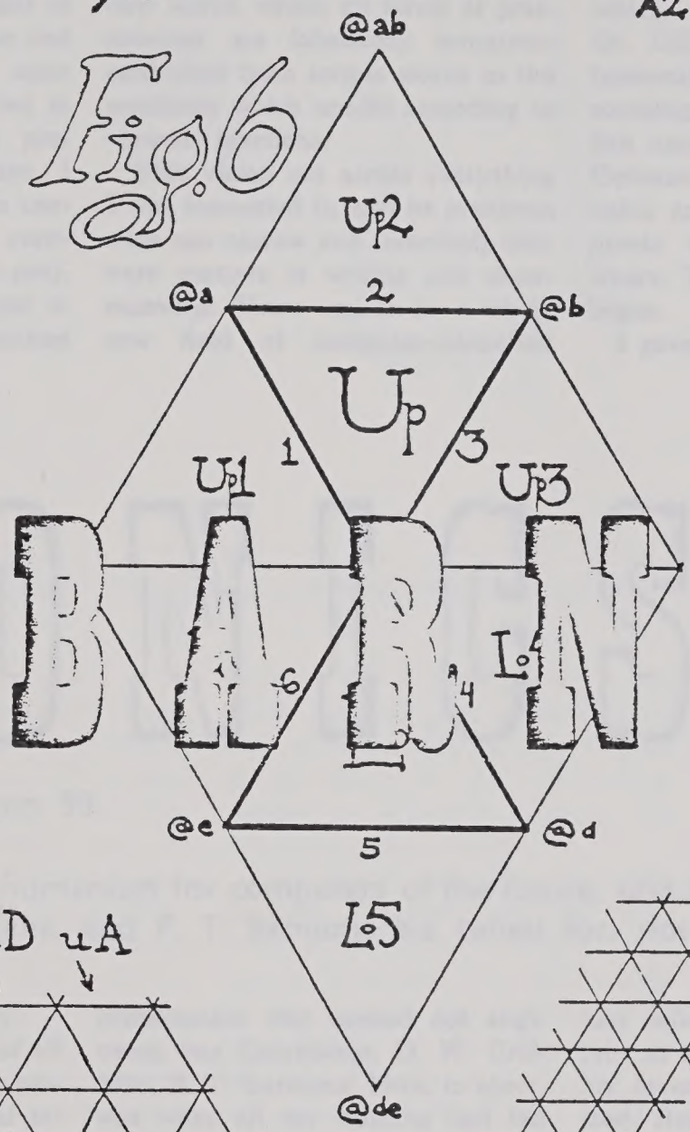

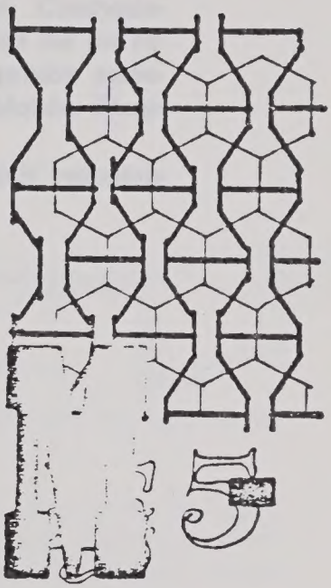

BARNUM

by Theodor H. Nelson '59

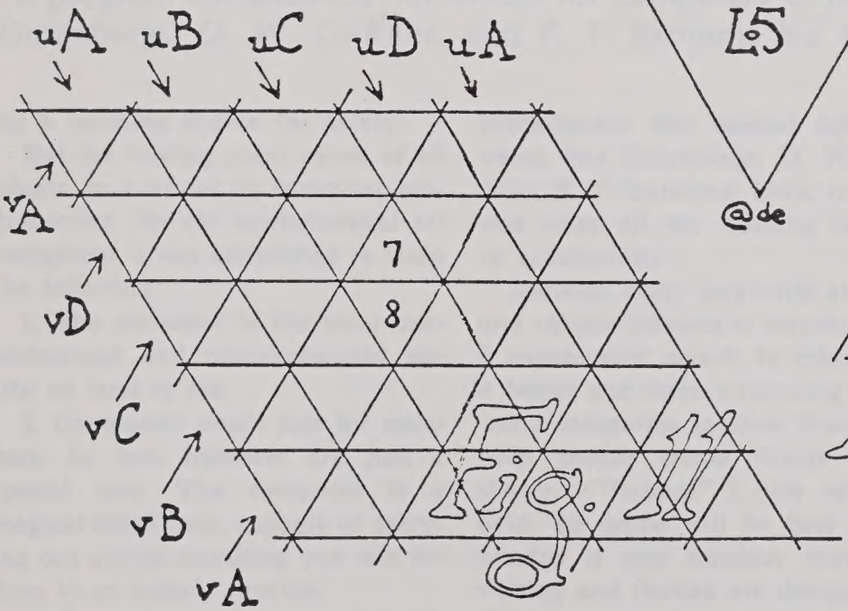

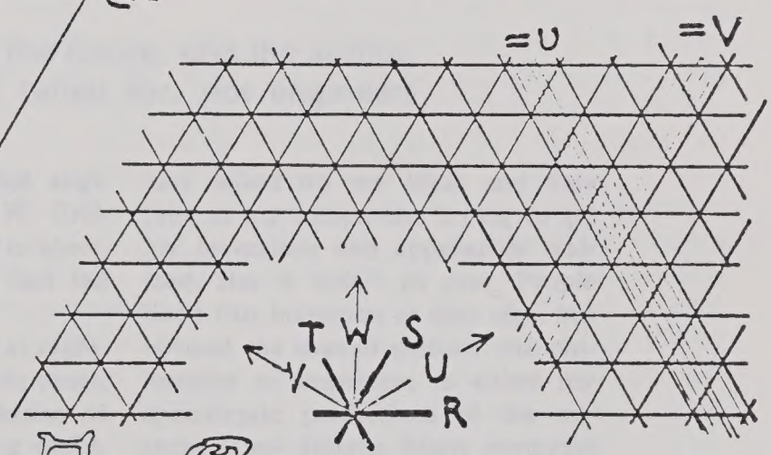

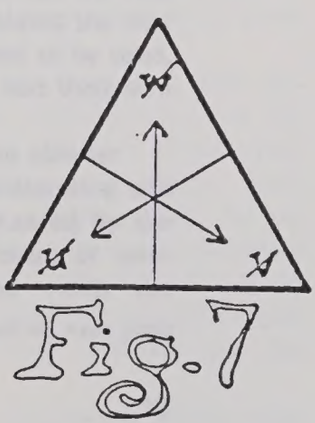

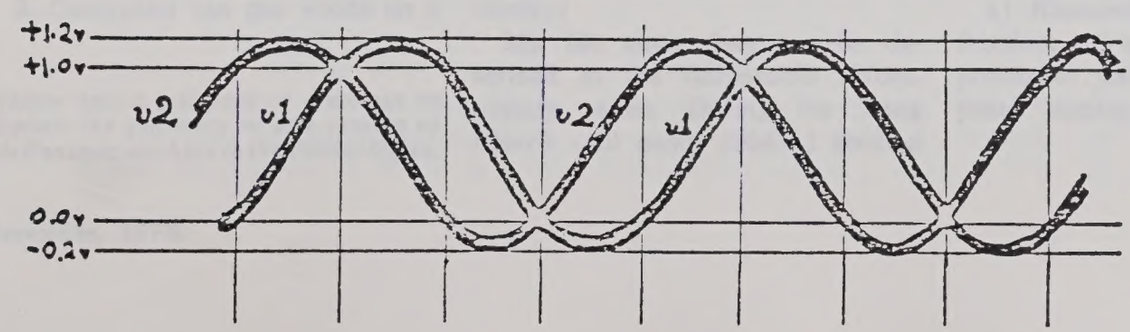

TRONICS

as told to Theodor H. Nelson '59

To program Renaissance humanism for computers of the future, says the author, Gutenbergs, D. W. Griffiths, and P. T. Barnums are called for, not engineers

A turning point in your life is not generally something you plan in advance. When I got out of Swarthmore I wanted to be a movie director, so what I was doing in graduate school is hard to explain. But I thought of myself as a writer and showman and looked back with enjoyment upon various innovations I had pulled in magazines and shows I had produced. Perhaps most important, I had developed an immunity to conventional advice, based on the eventual vindication of various large projects. I looked forward to a career in writing and films after I had picked up a teaching degree for safety.

But the turning point came, of all places, in a course on computer programming. An old mathematical incompetent, I was astonished to learn the following:

1. The computer is the most misunderstood and misrepresented entity on land or sea.

2. Computers aren't just for numbers. In fact, numbers are just a special case. The computer is a magical detail man, capable of carrying out almost anything you can reduce to an orderly process.

3. Computers can put words on a screen and respond to actions by somebody at a keyboard.

4. Computers can make pictures.

Suddenly it was all clear to me. There was soon going to be a whole new world, where all forms of presentation are fabulously computer-controlled from scripts stored in the machines which unfold according to viewers' reactions.

This vision cut across everything I was interested in, and its problems were not narrow and technical; they were matters of writing and showmanship! There was to be a whole new field of computer-controlled presentation that needed not engineers, but Gutenbergs, D. W. Griffiths, P. T. Barnums! Here, in short, was what all my training had led to accidentally.

Through many long walks at night, and various sessions of leaping heart, I consecrated myself to creation of a better and more interesting world, using computers to show things and help people create things to be shown. ("Fantics" I now call this field. Its scope will be best apprehended if you consider that both writing and movies are things to be shown.)

My life since then can be described in the fashionable revolutionary terms. During the Long March (till about 1966) I studied computers on my own with unending phone calls, attendance at conferences, and reading, especially the manufacturers' free literature. I made a living any way I could, which included a sojourn with Dr. Lilly's dolphins in Miami, professional folk singing, and teaching sociology at Vassar. I began to publish and speak at conferences. The Defense Ministries of Czechoslovakia and Norway asked me for reprints. I felt I was getting somewhere. The Third, or Mobile, Phase began.

I gave papers and made proposals and talked up my ideas and took jobs at big companies, trying to get my inventions and approaches realized. But it didn't go over. People liked this invention or that idea, but refused the overall picture, the philosophy so important to either my apocalyptic predictions or the nuances of my designs. Many computer people seemingly didn't like my stuff because they felt it violated the way God intended computers to be used. And laymen evidently had their own reasons for distrust.

Yet my message is so simple:

1) Knowledge, understanding and freedom can all be advanced by the promotion and deployment of computer display consoles (with the

continued on next page

Nelson uses a full page of drawings to explain the geometry of one version of his Fantasm machine to the Patent Office.

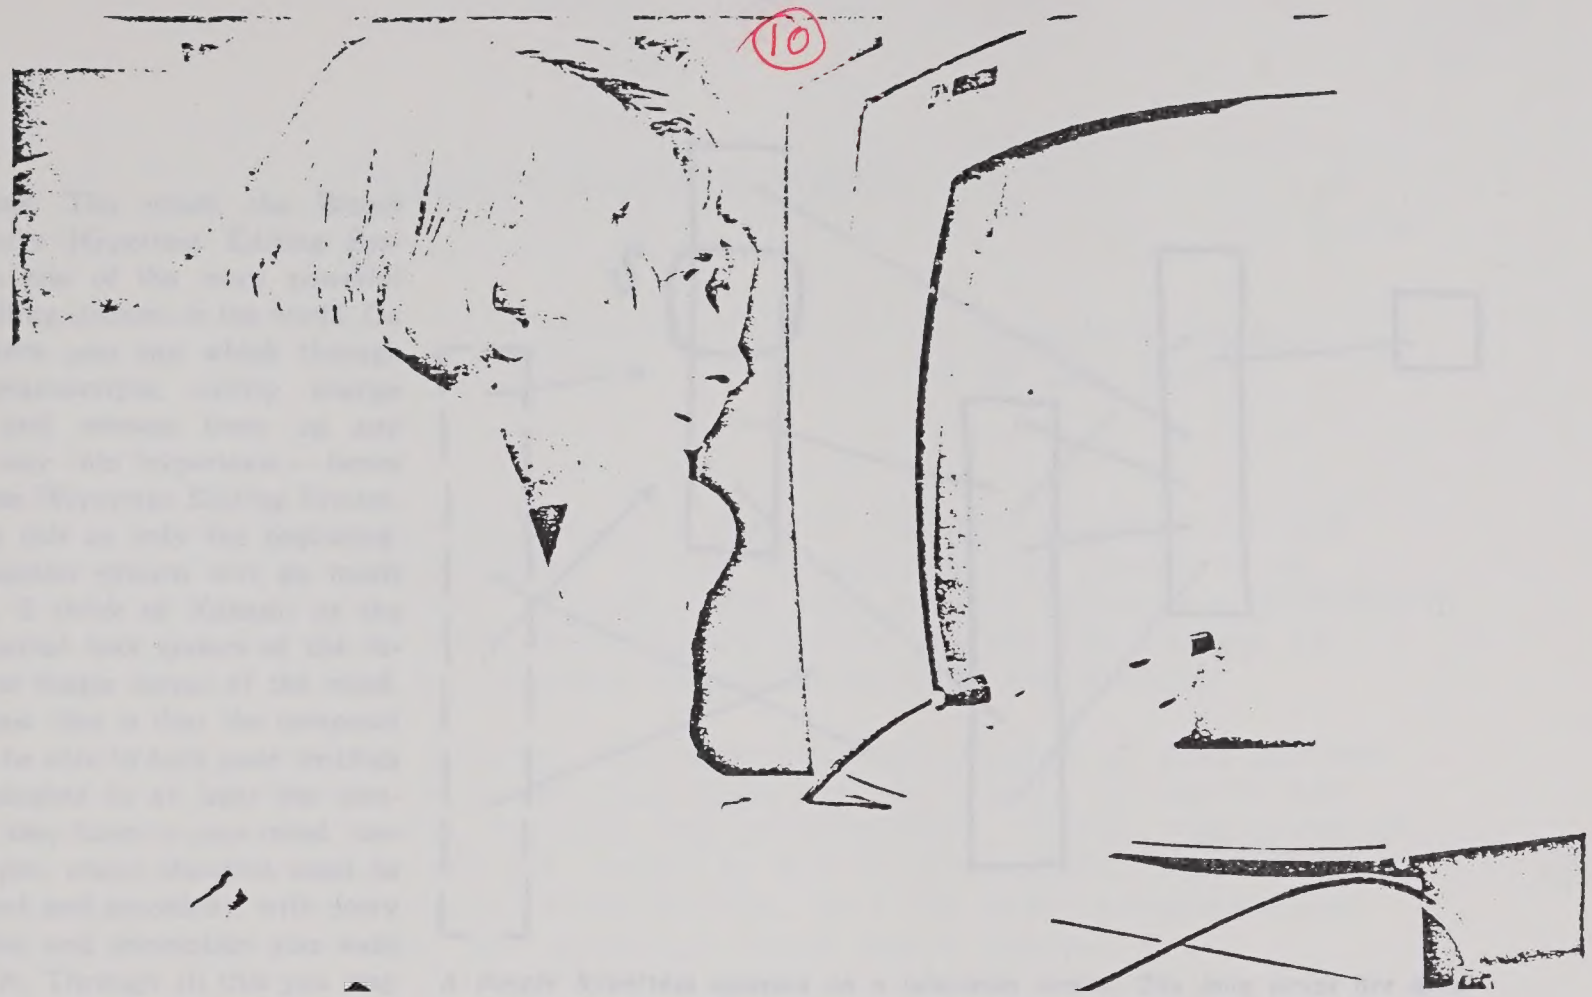

At the 2250 console, Inventor Nelson anticipates the forthcoming "age of prestidigitative presentation and publishing."

right programs behind them).

2) Computer presentational media, coming soon, will not be technically determined but rather will be new realms for human artistry. This point of view radically affects how we design man-machine systems of any kind, especially those for information retrieval, teaching, and general writing and reading. Some practitioners see such systems as narrowly technical, with the computer hoisting up little pieces of writing on some "scientific" basis and showing them to you one grunt at a time. A Metrecal banquet. I disagree. The systems should be opulent.

3) The problem in presentational systems of any kind is to make things look good, feel right, and come across clearly. The things that matter are the feel of the system, the user's state of mind, his possible confusion, boredom or enthusiasm, the problems of communicating *concepts*, and the very nature of concepts and their interconnection. There will never be a "science" of presentation, except as it relates to these things.

4) Not the nature of machines, but the nature of *ideas*, is what matters. It is incredibly hard to develop, organize and transmit ideas, and it always will be. But at least in the future we won't be booby-trapped by the nature of paper. We can design magic paper.

It is time to start using computers to hold information for the mind much as books have held this information in the past. Now information for the mind is very different from "information for the computer" as we have thought of it, hacked up and compressed into blocks. Instead we can stretch the computer.

I am proposing a curious kind of subversion. "Let us design," I say; and when people see the systems, everybody will want one. All I want to do is put Renaissance humanism in a multidimensional responsive console. And I am trying to work out the forms of writing of the future. Hypertexts.

Hypertexts: new forms of writing, appearing on computer screens, that will branch or perform at the reader's command. A hypertext is a non-sequential piece of writing; only the computer display makes it practical. Somewhere between a book, a TV show and a penny arcade, the hypertext can be a vast tapestry of information, all in plain English (spiced with a few magic tricks on the screen), which the reader may attack and play for the things he wants, branching and jumping on the screen, using simple controls as if he were driving a car. There can be specialized subparts for specialized interests, instant availability of relevancies in all directions, footnotes that are books themselves. Hypertexts will be so much better than ordinary writing that the printed word will wither away. *Real writing by people*, make no mistake, not data banks, robot summaries or other clank. A person is writing to other people, just as before, but on magical paper he can cut up and tie in knots and fly around on.

A few of my ideas have been put into practice. Andries van Dam '60, now associate professor of computer science at Brown University, instigated a text project partly at my urging. Taking off from a document I wrote, he and his students put together a big computer program which we argued about endlessly over the

telephone. The result, the Brown University Hypertext Editing System, is one of the more powerful text-editing systems in the world. On the screen you can whisk through your manuscripts, swiftly change them, and connect them up any which way into hypertexts — hence the name Hypertext Editing System.

I see this as only the beginning. My Xanadu system will go much farther. I think of Xanadu as the fundamental text system of the future, the magic carpet of the mind. The basic idea is that the computer should be able to hold your writings and thoughts in at *least* the complexity they have in your mind (unlike paper, where thoughts must be truncated and parodied), with every cross-link and annotation you want to put in. Through all this you may zoom like a bird in an enchanted forest. The system will help you ponder complex theories and compare variations of what you're studying or creating; it should also allow you to go back in time through earlier versions of your work, perhaps building again on drafts you thought you had discarded. You can sift and combine your notes into a conventional work or leave it all hanging in a huge controlled agglomeration. The system will help integrate syntheses, unravel inspirations, deconfuse thought. But, of course, you may read and write hypertexts. Every kind of human creativity — not just writing—can be aided if we build a sufficiently general creativity console, such as Xanadu.

Although early systems of this type will cost unspeakable amounts of money, later in the seventies it should be possible to outfit an entire college campus, for example, for a few thousand dollars per console. Think of not having to hand in your seminar paper physically; zip it instead to the antechambers of your readers' consoles with the bump of a button.

Besides these visions, which only get technical at certain key points,

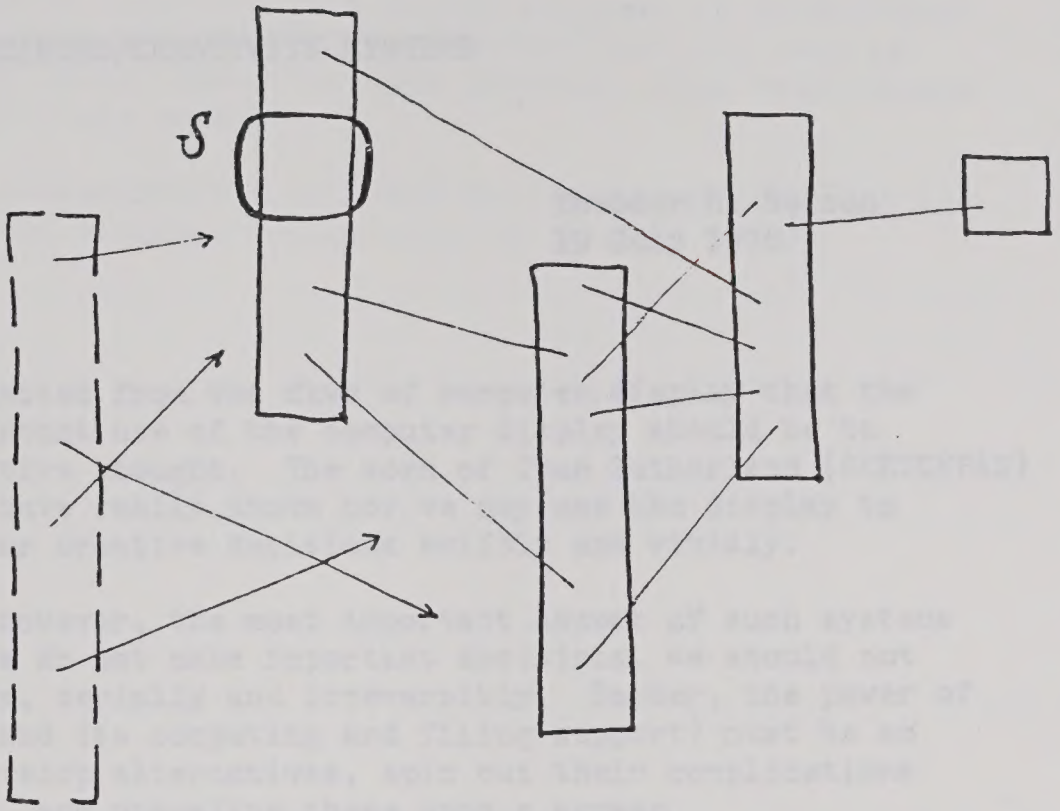

A simple hypertext appears on a television screen. The long strips are documents it can roam over. Not only can viewer leap along the connection lines indicated; he can also add indexes and commentaries to help in mutual comparison (dashed strip)—a facility of Xanadu system, "the magic carpet of the mind."

I have also worked all these ten jobs (would you suppose my hypertexts were relevant to the ABM system? Would you believe the telephone company?). Until it's time.

So far my predictions have been generally right except for chronology. I originally thought the printed word might be eliminated by 1970 or 1972. Now, uh, I guess it will take a little longer. ("Is Nelson paranoid?" asks a recent letter to *Computer Decisions* magazine.) But it's going to happen. Computer screens will be in the home, perhaps sooner than in the school. No more graveyards of paper for the words we write. No more pencils, no more books, no more teachers' dirty looks.

Ladies and gentlemen, the age of prestidigitative presentation and publishing is about to begin. Palpitating presentations, screen-scribbled, will dance to your desire, making manifest the many mysteries of winding wisdom. But if we are to rehumanize an increasingly brutal and disagreeable world, we must step up our efforts. And we must hurry. Hurry. Step right up.

(12)

DCQ 1
b1

DECISION/CREATIVITY SYSTEMS

Theodor H. Nelson
19 July 1970

It has been recognized from the dawn of computer display that the grandest and most important use of the computer display should be to aid decisions and creative thought. The work of Ivan Sutherland (SKETCHPAD) and Douglas Engelbart have really shown how we may use the display to visualize and effect our creative decisions swiftly and vividly.

For some reason, however, the most important aspect of such systems has been neglected. We do not make important decisions, we should not make delicate decisions, serially and irreversibly. Rather, the power of the computer display (and its computing and filing support) must be so crafted that we may develop alternatives, spin out their complications and interrelationships, and visualize these upon a screen.

No system could do this for us automatically. What design and programming can create, however, is a facility that will allow us to list, sketch, link and annotate the complexities we seek to understand, then present "views" of the complexities in many different forms. Studying these views, annotating and refining, we can reach the final designs and decisions with much more in mind than we could otherwise hold together in the imagination.

Some of the facilities that such systems must have include the following:

Annotations to anything, to any remove.

Alternatives of decision, design, writing, theory.

Unlinked or irregular pieces, hanging as the user wishes.

Multicoupling, or complex linkage, between alternatives, annotations or whatever.

Historical filing of the user's actions, including each addition and modification, and possibly the viewing actions that preceded them.

Frozen moments and versions, which the user may hold as memorable for his thinking.

Evolutionary coupling, where the correspondences between evolving versions are automatically maintained, and their differences or relations easily annotated.

In addition, designs for screen "views", the motion, appearance and disappearance of elements, require considerable thought and imagination.

The object is not to burden the user, or make him aware of complex-

ities in which he has no interest. But almost everyone in intellectual and decision pursuits has at some time an implicit need for some of these facilities. If people knew they were possible, they would demand them. It is time for their creation.

A full-fledged decision/creativity system, embracing both text and graphics, is one of the ultimate design goals of Project XANADU.

WORKSPACES, TEXT THEATERS, AND HYPERTEXTS: PURPOSE OF THESE NOTES
Theodor H. Nelson

Hypertext Note 0 [29 Apr 1967]
Private circulation only. This will eventually be incorporated in a longer article. Terminology is tentative.

Interest in hypertexts leads directly to a concern with text theaters. You cannot make or present a hypertext without a text theater; and the less limited you want the hypertext to be, or the less sure you are of what your final product will look like, the more versatility you need in your theater.

Likewise, unless one wants his hypertexts to conform to fixed and simple rules, and appear in fixed and simple formats and activity spaces, then the nature and creation of activity spaces becomes of great concern.

Thus there are several things one can do about hypertexts.

1) One can try to create hypertexts without a theater, which is difficult and probably pointless.

2) One can endeavor to figure out in the abstract what kinds of hypertext (and activity space) are possible, and their properties.

3) One can design text theaters in which the creation and use of hypertexts (and activity spaces) is possible and practical.

These hypertext notes report principally on work of types 2 and 3. Because of the interest in this work among those in computer text handling, these notes are intended to short-circuit the publication cycle, and get various ideas abroad expeditiously.

(15)
(2)

BRIEF WORDS ON THE HYPERTEXT
Theodor H. Nelson

Hypertext note #1. 23 Jan 67
Private circulation only. This will eventually be incorporated in a longer article.

About the term. 'Hypertext' is a recent coinage. 'Hyper-' is used in the mathematical sense of extension and generality (as in 'hyperspace,' 'hypercube') rather than the medical sense of 'excessive' ('hyperactivity'). There is no implication about size-- a hypertext could contain only 500 words or so. 'Hyper-' refers to structure and not size.

What kinds of structure are possible in a hypertext? Any. Ordinary text may be regarded as a special case-- the simple and familiar case-- of hypertext, just as three-dimensional space and the ordinary cube are the simple and familiar special cases of hyperspace and the hypercube.

(It should be understood that all other examples of hypertext will also be special cases. There is no 'true' form of hypertext, only a very large number of possible structures.)

Ordinary text, called by McLuhan and others "linear," is a continuing sequence in a fixed order. (Actually there are many departures from this: the footnote is the most obvious, but the summary, illustration, headline and other features also break away.)

But from the standpoints of both subject matter and presentation, it is rarely ideal to have a single fixed sequence of materials or ideas. And there are ways now to present useful text structures of every conceivable kind.

To generalize the idea of ordinary text, we need to find a means of representing departures from mere sequential order. This lies in the mathematical theory of graphs. A graph structure is a web of connections between points. It may have connectors that have directions ("arcs")

or it may have undirected connectors ("edges")

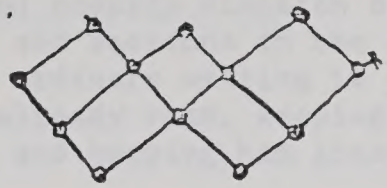

and it may even have arcs or edges of different kinds ("colors"):

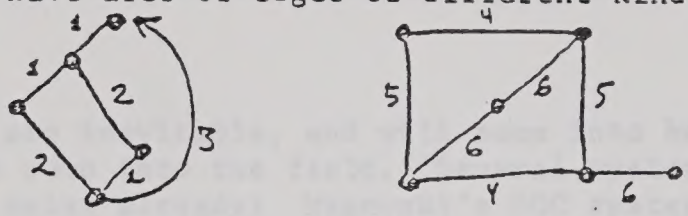

Arcs, edges and colored connectors all have possible analogues in the hypertext.

Each dot represents a chunk of text. The lines represent connections between chunks of text. A plain line, or edge, represents a connection of possible travel from one chunk to another. A directional line, or arc, represents a connection of possible travel from one chunk to another in <u>one direction only</u>.

Ordinary text is like this:

However, since it is possible to backtrack or skip, it is really more like this:

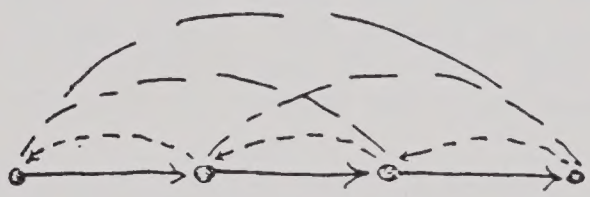

In this case, the arcs represent prefabricated or convenient connections (directional), and the dotted lines represent optional travel (in either direction).

Finally, "colored" arcs and edges-- that is, arcs and edges that have been assigned to different types-- may represent <u>conditional</u> travel. For instance, the author/editor of a hypertext might arrange that only persons having certain qualifications (or persons who had covered such-and-such material) should be allowed to pass along a certain connection. The numbered (or 'colored') connector types may represent these permissions, and other conditionals.

Hypertext diverges from ordinary text in that the reader's possible sequences diverge from plain sequence. Is he to have choices? And how are they to be expressed?

The answer is that the reader has the choices that the author/editor allows him. He may not be aware that choices are made, and may think he is reading a book and expressing his opinion of it. Or at the other extreme, he could be given complete and continuing information on where he is in the overall text.

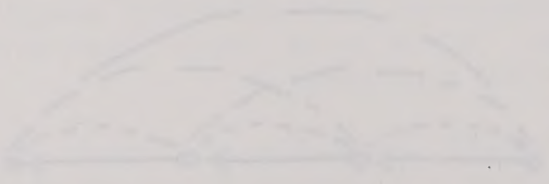

In ordinary writing there are two problems. The first is structure: creating a suitable and orderly division of the subject, and creating appropriate sequences and sections in the text for its proper presentation. The second problem of ordinary writing is presentation: keeping the reader aware of what he has already read, keeping him oriented as he sees different parts of the subject, and keeping him interested.

These problems are unchanged in the hypertext.

Hypertexts are inevitable, and will come into being just as fast as text handling systems come into the field. Several systems suited to handling some hypertexts exist already: Magnuski's DOC system at Project MAC, Bernstein and Slojkowski's program documentation system at Bell Labs, and Engelbart's editorial console at Stamford Research Institute.

(18)

HYPERTEXTS and, or perhaps versus, CAI
Theodor H. Nelson

Hypertext note #2 23 Jan 67
Private circulation only. This will eventually be incorporated in a longer article.

Programmed teaching has been improved by the introduction of computers. These make it possible for a teaching program to branch, and follow a graph structure. (Thus it may, of course, be a type of hypertext.)

Practitioners of this art-- the stringing of program "frames" for teaching-- are generally of the behavioristic persuasion, believing a) that it makes sense to define the goals of teaching simply in terms of the subject's resulting behavior, and b) that a construal of the teaching task in terms of reinforcement schedules and evoked behavior successfully comprehends the whole problem. To question these assumptions is pointless, since they are either self-evidently false or there is no point in arguing.

Anyway, it is believed in this field that since the problem is one of "learning" (in the pigeon/rat sense), the proper way to teach is by a highly detailed succession of presentations and fill-ins, minutely crafted to provide explanation and "reinforcement" in precise and tiny doses.

The alternative viewpoint I would like to offer is that the problem of teaching is motivation (and comprehension, which raises or lowers motivation) and the way to teach is to present things so as to maintain high interest. It is necessary to permit individual, personal success in an individualized study environment-- where the student makes a field his own, immersing himself in those aspects of the field that appeal to him, probably without the interfering personification (and perhaps gratuitous interference) of the teacher. Understanding, autonomy, success at exercises, and motivation must all grow together.

Large hypertexts, with many divergent but cross-cutting regions of study and activity, seem to be the way to do this.

Without harping, let me list several salient distinctions between this hypertext approach to education and the CAI approach.

1. Reinforcement vs. Motivation. A motivated individual will arrange to be reinforced. A reinforced individual will not necessarily arrange to be motivated.

2. Minutiae vs. sweep. CAI requires infinitesimal attention to particulars. Hypertexts require considerable attention to generalities.

3. Memorization vs. Comprehension. The old sense of comprehend was "comprise"-- press together a variety to form a whole. (This is the whole point of survey courses and outside reading; yet by making these compulsory the most important factor, selective personalization, is lost.) Memorizing is not understanding.

4. Acceptance vs. Considering Alternative Viewpoints. All fields have areas of controversy, or at least (as in mathematics) entirely different ways they can be considered. Rather than enforce the initial acceptance of a particular viewpoint-- which is why children disaffect from subjects and from school-- we must give variety and alternatives. What we like best we learn best; and this is no dead end, but a framework for the rest.

5. "Learning the basics" vs. Giving Alternative Handlings and Dismantlings. The myth of education is that there are "basics" in every field. This is simply not true. There are indefinitely many ways that each so-called subject can be dismantled for teaching. Good teaching provides handles and approaches for a sense of understanding, and alternative sets of "basics" in the hope that one will appeal.

6. Difficulties of Production: Crossword Puzzle vs. The Panorama. The difficulty of making good teaching programs is notorious. As a form of activity it resembles nothing more than creating crossword puzzles. Will it be like this to write teaching hypertexts? Probably these will be difficult, but not so compacted and intricate. The central thing we may depend on is this: someone who knows what he is reading or doing, and has freedom of movement, can be relied on to get himself out of a meaningless dead end. If he is being led blindfolded, without explanation, through a maze of responses and reinforcements, everything must be correct; for if he is stalled he cannot help himself, and he cannot help you.

Another possible analogy is the difference between making animated cartoons and epic films. They both take money and trouble. One requires great attention to specifics of detail in interlocking fragments; the other requires equipment and material on a grand scale, and logistic and aesthetic direction.

We may think of CAI as an extension of programmed teaching, of the general hypertext as a generalization of writing and editing. Whether findings from CAI will have a bearing upon hypertexts for teaching may be doubted.

A word on Socratic CAI. "Socratic" systems have a touch of humanism in them, as they try to express conversationally the exact understandings and misunderstandings of the pupil. This is a wonderful technique, but is it universal in application? Consider: would you rather have directions to get to a place, or a road map? The former introduces spurious social interaction, an unnecessary problem in most cases, and does not provide for errors either by the director or the follower. A road map allows self-correction.

Moreover, the Socratic system has a tendency, at least in the published examples, subtly to put down the student while it promulgates the myth of the single right answer. This is what we are doing in education already, to the national detriment.

(7)

TEXT AND TEXT-RELATED DISPLAY: SOME TERMS
Theodor H. Nelson

Hypertext Note 3 7 April 67
Private circulation only. This will eventually be incorporated in a longer article.

There will be many ways of displaying texts on computer display screens, and many reasons for doing so. There will therefore be different kinds of theaters, or settings, for handling texts upon a screen; and many kinds of parts and arrangements in which the texts will appear. In this and later notes various types of setting, and different parts of displays, will be distinguished. I hope these terms will make it easier to describe, and think about, present and future systems. The overall intent is to clarify a lot of needed ideas, and assign the most appropriate terms to them.

Theater. Let us use the term "theater" (or perhaps "setting") for a display console that has various capabilities for displaying, filing and computing. Display system has been in common use, but this term has problems. "System" could refer to the console alone, the program and file facilities available through it, or the overall computer environment-- time-sharing, monitored or whatever-- in which it functions. It is simpler to speak of a "theater," restricting that term to the console and the facilities available through it. By this terminology, then, one system can have many theaters -- text theaters, engineering design theaters, and so on.

After all, the dictionary sense of "theater" is "a place where certain things occur." (Webster III, definition 4a: "a place or sphere that is the scene of dramatic events or significant action.") (I think this term "theater" is also preferable to "facility." "Facility" best refers to some acting part that does a specific kind of job; for which the term "capability" is not really satisfactory, and for which the term "capacity" should likewise not be used-- the capacity of the system being how much it can hold, in its various buffers, buckets, cisterns and oceans.)

A facility, then, is an acting part of the system or theater that does a specific job. Possible text facilities include the filing and retrieval of texts, word and phrase searching, graph following and graph structure analysis, the placing of markers and connectors in the text, and, of course, simple text editing.

An information base is the information on which displays are founded. Modifying the information base may take place in some systems, not others. Text editing is one way to modify an information base; rebuilding a SKETCHPAD structure is another. Different systems will have different facilities for modifying the information base.

The file structure of a system is the way in which information is divided, connected and stored. The extent to which a file structure will influence the material stored on the system will vary. Many features will in part determine how much the user is affected by the file structure. These will include types of connectors, record length, restrictions on connectors, categorization methods and many other aspects.

(21)

WORK THEATERS, WORKSPACES AND THEIR PARTS (8)
Theodor H. Nelson

Hypertext Note 4 7 April 67
Private circulation only. This will eventually be incorporated in a longer article.

A spread (screenful? presentation? display? frame? shot? event?) is a collection of text and graphics which is output as a complete screen display.

A window is an area on a screen which "looks into" a connected body of text or other information. However, it does not add to or rearrange material found there. If the text can be made to move through the window ("scrolling"), we may also think of this as the window looking at successive parts of the text file. Thus a placemarker or window pointer (or sashweight?) is a pointer into the text file by which the window's location in that file is maintained.

This "window" idea is useful mainly for free text or lists. For coded information--legal, survey, municipal, bookkeeping-- forms are required. A form is a pattern of words or information on the screen to which is added information obtained from a file. A window that is empty, that is, pointing to an empty area, is an unoccupied window.

A fixed window always occupies the same place on the screen. A movable window may be moved about on the screen by the user (either continuously or in jumps.) A dependent window will take its position from the position of something else on the screen. An optional window may be made to appear or not, at the user's preference. Alternative windows are sets of windows of which only certain windows will appear at a given time.

A layout or format (short for screen layout or screen format) is an arrangement on the screen of different windows, forms, system messages, and choices. (However, system messages and alternatives may not necessarily appear on the screen.) Alternative and conditional formats may be useful.

A sublayout (subformat) is a part of a format (or layout) which may itself have alternatives, be replaced, move around, or appear in different formats.

A free message or loose message is a text message not in a window or form.

A face (door? portal? wall? side?) is a format assigned to a position in an overall workspace.

A workspace or work structure (hyper-workspace? display space? implicit conceptual topology?) is a structure of formats, arranged, in a branching graph, so as to succeed one another upon the screen. A graph of the connections of these formats we may call an inter-format graph. Through these formats the user may choose information and displays, negotiate with and modify the information base, call facilities and enter new information, as required. The construction of suitable workspaces for specific tasks will be a difficult and complicated matter, with personal preferences playing a major part.

Steps from format to format in the workspace we may call jumps, steps or perhaps cuts (as in the movies). Conditional, optional and alternative formats might be desirable at specific vertices.

We might use the term hyper-desk for a console at which all your work can be done, or a significant mix of tasks.

The access structure (or access space) of a workspace is a graph of the ways in which specific information may be reached from different layouts. This may not have the same graph structure as the workspace itself, since the same information may presumably often be reached from different formats.

CREATIVITY SYSTEMS
Theodor H. Nelson

Hypertext Note 5 7 April 67
Private circulation only. This will eventually be incorporated in a longer article.

Creativity support system (or creativity system) is a term I wish to reserve for a particular kind of text-handling theater and file system.

To explain this, a word is necessary on the creative process.

Constructive creativity-- the kind that occurs in writing, theoretical thinking and design-- has the following stages. A number of parts, or entities, are seen to hang together under certain ideas.

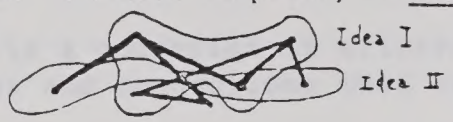

The entities are then arranged or structured according to the ideas. (For instance, in writing, a number of points are put together under an outline.) Typically, however, these things do not fit well together at first, and various combinations and arrangements have to be tried out and compared. Ordinarily we can only try a few arrangements, often only one, which is patched, rearranged and repaired until we have to stop. The ideal creativity support system, then, would help us put together, speedily and easily, all the different arrangements we wished to compare, and let us compare them quickly and conveniently.

The following, then, are the criteria of a creativity support system:

1) The system can hold alternate structures of the same materials, and accept new ones easily. If it is a text system, these materials will be texts; if a system for mechanical design, they might be alternative designs of a machine. (Alternative criterion.)

2) The user may easily compare the corresponding parts of these alternative structures. (Intercomparison criterion.)

3) The user may easily label and annotate these alternatives, in various ways and "levels," so that he can summarize and keep track of his increasing insights about them. (Annotation criterion.)

4) The file categories, label types, etc. must be easily changeable to what the user wants. (Protean criterion.)

5) It must be possible to accommodate entries of any length, or at least very long entries and very short ones.

We are now ready to discuss what I will call the general creativity console.

Such a console will be a creativity system, as described above. Moreover, it must be "general" in that it will lend itself not merely to text writing, but to musical composition, graphic arts, motion picture editing and so on. This does not necessarily mean it can show movies or perform musical compositions. But it must be able to hold written information about these things, and connect usefully to equipment that can show them.

This means that it must be readily hooked up to A) receive signals from other equipment, and B) control other equipment in turn. C) Its visual displays must be easily supplementable with new console button arrangements. (In hardware, for speed.)

More importantly, D) its operations must be easily changeable to new ways of working-- by a simple, noncomputerish command structure, perhaps a pictorial one. And finally, E) the record structure by which things are filed must be prevented from affecting the user: whether storage on disk or tape is by ten-character chunks or thousand-character pages, the user must be shielded from being aware of how these relate to the things he is doing. Otherwise the work may be influenced improperly by the distraction of such structures.

(23)

PRESENTATIONAL THEATERS FOR INTERACTIVE MEDIA
Theodor H. Nelson

Hypertext Note 6 7 April 67
Private circulation only. This will eventually be incorporated in a longer article.

Computer display makes possible a number of interactive media-- information media showing prepared materials and allowing the user various choices of what he will see. These include the hypertext, hypergram, annotated tableau and hyper-comic.

1. Texts and Hypertexts.

A text-- in the ordinary sense-- is a congeries of written words. These words are arranged in divisions and sequences and connections that seemed to the author or editor appropriate to the subject.

A hypertext, therefore, is similarly a congeries of written words, similarly arranged in divisions and sequences and connections that seem to an author or editor appropriate to the subject. But the hypertext is not restricted to a simple succession, and may therefore branch in different directions. Thus a hypertext is text arranged in a graph structure, with the branches to be made according to the choice of the user or the criteria set down by an author or editor.

The term "hypertext," as used above, can include anything from through-composed text structures to text structures which have merely been collected: as long as they are composed mainly of text sections and branches.

Presumably hypertexts may contain graphics of various kinds, described below. These might be called "illustrated hypertexts," or we might not bother to make the distinction.

2. Graphics.

A "graphic," by convention, is some sort of picture, symbol, design, or other intentional part of a visual display that is not made up of letters. We may distinguish among several kinds of graphics for computer display.

A literal graphic is a graphic which is pulled out of memory and put directly on the screen without variation or movement.

A sketch is a graphic which is modified or translated by the machine, adapting or abstracting from some data, but where the data was created for or in part resembles the final result.

A generated graphic is one which is created by the system from data, mainly ad hoc. The elements of its structure are found rather than prepared beforehand.

Moving or movable graphics are those which are capable of internal movements (as opposed to creeping or relocatable graphics, which simply occupy successive positions).

A hypergram or hypergraph is a graphic, or diagram, that can be manipulated, either freely by the viewer or by the system on the basis of data or other occurrences.

A tableau is a spread which is largely pictorial or graphical, but contains inset text matter in the form of captions, dialog, etc. (If the subjects are people and dialog, "hyper-comics" will perhaps be the better term.)

Such devices may seem far-fetched. However, their use for teaching the undermotivated and the poor reader may be considerable. And their possible use for more serious tasks is not unthinkable. According to one reporter, the President's guidebook to our nuclear warfare command codes is a sort of garish comic book, "pages of close text enlivened by gaudy color cartoons. They looked like comic books-- horror comics, really, because they had been carefully designed so that ... military aides could quickly tell him how many million casualties would result from Retaliation Able, Retaliation Baker, Retaliation Charlie, etc...."*

---

*William Manchester, The Death of a President, quoted in The New York Times, March 29, 1967, p. 40.

These terms may seem unnecessary. But they may help in simplifying the description of existing and possible display systems.

For instance, SKETCHPAD is a graphics drawing facility which permits the modification of a graph-structured data base by screen manipulations. It creates sketches from this data base, and will permit the creation of hypergrams in two and three dimensions.

Or, a picture of a machine or human body which could be "opened" and "manipulated" on the screen, with expandable labels, would be a hypergram tableau.

(25)
(12)

HYPERTEXT STRUCTURES
Theodor H. Nelson

Hypertext Note 7 26 April 67
Private circulation only. This will eventually be incorporated in a longer article.
Terminology is tentative.

There are a number of different types of hypertext. They differ principally according to the freedom of the user, whether and to what extent they are hierarchically structured, whether they are continuous or in discrete parts, and the way they are addressed by the user.

Their editorial qualities, and editorial techniques that may be useful, will not be described here.

(It would be convenient if we considered all features of hypertext to be "structural," a matter of graph connections. This is not the case. A number of interesting properties of hypertext are not structural; and hypertexts with 'the same' structure may be presented in ways that are importantly different to the user. The way the system faces the user is rather independent of the connective structure of the text.)

A text is not a hypertext if interaction is irrelevant and the user's actions have no effect on what is seen. If what you see depends on qualities you bring with you, what you already know, or what you are, then it should be called predetermined conditional text, or "fated" text. (However, a hypertext might be partly fated.)

Freedom and explicit control by the user are important attributes. If the user is given orientation and explicit ability to move freely in the text, it is free hypertext. If his movements determine what he sees in ways he cannot control, it is a conditional hypertext, or text maze. If the user is not given overall or orienting information about the text graph he is in, he does not maintain control. Though his actions may determine what he sees, if he doesn't know it and can only guess at the structures with which he is communicating, it cannot be called 'control.' (Conventional computer-assisted instruction is almost exclusively of the text maze type.) A hypertext may, of course, be partly conditional and partly free.)

(A presentation or structure may also be available depending on some action by the user which occurred earlier but not just before the current one. In this case we may call it a back-conditional.)

A hypertext may be hierarchical or non-hierarchical. There is a strong bias among technical people (fostered perhaps by the tradition of the Harvard outline of topics: I.A.1.a. ...) to suppose that hierarchy is natural and universal. This is a complex assumption about the nature of verbal information and ideas. Hierarchical hypertext forms include paragraph trees, stretchtext, etc. But non-hierarchical hypertext forms are possible which could take on any graph structure whatever. How this will usefully relate to content and function, and how the user will keep track of where he is, are problems.

Nevertheless, certain types of non-hierarchical connections should definitely be useful in hypertext. These include "return to start," and material set aside for modal queries (see below).

Continuous and discrete structures are both possible in hypertext, as they are in ordinary text. (By "continuous" we mean something like "smoothly written.") This is partly an issue of whether presentational breaks in the text are formally known to the system (like, in print, the paragraph, chapter break and boldface heading) or simply matters of editorial character.

(The extent to which this continuous/discrete distinction will be part of the computer software and technical properties of hypertexts is not clear. For instance, stretchtext seems intrinsically continuous, but users could elect to give it discontinuities that might be useful.)

The possibility of continuous hypertext with other than a simple sequential structure should not be overlooked. Weird topologies might be useful and possible in some unthought-of way.

(26)
(13)

The way in which choices and branches are presented to the user may vary independently of the connective structure by which they are strung together. There may be choice points (or breakpoints), where the presentation (or text) stops till the user chooses; or the choices may be imbedded in branch points-- markers in a continuing presentation that indicate the user's branching options. (If we consider that a knothole on a wooden board means a branch once went from there, knothole might be a xylologically appropriate term for branchpoint markers.)

It may be that branch possibilities are not shown at all. (Announced vs. unannounced links.) They might be forbidden to the user, which is probably a poor idea. Or he might have to guess what options are available, by attempting queries. Whether this would serve a function is the reader's guess.

The reader might have the option of suppressing branch-option signals in order to concentrate, and reviving them when he wanted to digress. This would be another matter.

We may distinguish between standing options (always there) and present options (that come and go). There may be modal branch options, different types of branch to connected texts. If all types of modal option are always available, they are standing modal branch options. (Standing modal options could make it unnecessary to signal choicepoint and branchpoint links; there would always be n connections, or perhaps m points of curiosity the user could have satisfied at any time.)

It should be noted that modal branch options are surely not hierarchical, but partake of some other graph structure.

Branches or jumps may occur from any point in text to any other point in text. Ordinary writing usually permits jumps from the end of a sentence or clause to the beginning of a sentence (the footnote); but there are other possibilities.

Branches may occur between two distinct units (like the footnote), or between two places within the same unit, in which case it should be called a jump or hop. (In conventional text they say, "see page 3.") These jumps may be bidirectional or one-way.

(An esoteric question about branching is associated with hierarchical hypertext, such as stretchtext. The question is whether a jump within the same unit should also be to a point at the same level; and whether the branchpoint (e.g. in stretchtext) should be continuously available into further levels down.)

Questions, instead of explicit choices, may be used to manage user response at choicepoints. But this raises interesting issues. Questions may be topologically and functionally identical to free choicepoints. But are they psychologically?

If the questions give no hint as to what the answer will make the system do, then of course their effect is very different, for we have a text maze. But if the user is able to divine the probable result of the answers-- and is thereby motivated to guess, lie, etc., as a way to express choices-- the situation is complicated. Assuming that he is not being judged, graded or punished-- conditions whose psychological effect is not now evident, but it probably won't be good-- there is still that spurious conversational confrontation whose virtues may be less than supposed. (See "conversational systems," another note.)

The connective texture of hypertexts may be of many kinds, following any discrete topology (and perhaps several continuous). The connective texture may follow a regular pattern, like the graphs of mathematical groups; this would be a sufficient condition for creating a standing modal query structure. Or the connective texture of hypertext may be locally variable. The advantages of this have yet to be explored.

There may be one text type (as in most books) or several, modal types, distinguished in their meaning and presentation (like the captions, subheadlines and inset boxes in a magazine.)

Variants of areas or units or sentences may be presented to different users for different reasons; it may not make sense to speak of these as truly alternative units. There is no convenient way to discuss them here.

Randomness of system response might also seem useful, for enlivenment of dull material, challenge, etc. But it might also be demoralizing. The effects of randomness and other de-structuring is not of current interest.

(14) (27)

STRETCHTEXT
Theodor H. Nelson

Hypertext Note 8. 29 Apr 67
Private circulation only. This will eventually be incorporated in a longer article. Terminology is tentative.

Stretchtext is an extremely simple but powerful form of hypertext. It is probably the easiest possible hypertext to understand. It would be hard for a reader of stretchtext to become confused.

Stretchtext consists of ordinary continuous text which can be "stretched" or made longer and more detailed. By pointing at specific areas and pulling the throttle in the "magnify" direction, the reader may obtain a greater detail on a specific subject, or area of the text. The text stretches, becoming longer, with replacement phrases, new details and additional clauses popping into place.

The good of this structure should be evident. The reader remains oriented. If he loses track of where he is, he "shrinks" the text to a higher, smorter level; if he wants to study a topic in more detail, he magnifies it.

An important editorial constraint on stretchtext, then, is that details and narrative arrangements must remain fixed in their relative order through different levels of stretchtext. However, in one respect it appears to be easier to write than ordinary text: rather than deciding what details to "put in" and "leave out," the author merely assigns altitudes (or "fineness"?) to topics and details, thus determining at how great a magnification they will be seen.

High-performance dynamic consoles may be preferable for stretchtext, but ways are possible to present it on a static display-- if the display can support a moving cursor and a user pointing device.

Editorially, the stretchtext is (1) always the same unit, and (2) always a continuous narrative. Thus it is unlike hypertexts with discrete chunks and breaks.

A HYPERTEXT STRUCTURE FOR SELF-TEACHING
Theodor H. Nelson

Hypertext Note 9. 29 Apr 67
Private circulation only. This will eventually be incorporated in a longer article.
Terminology is tentative.

The best teaching is self-teaching. The hypertext design described in this paper is intended to be simple and clear, and help the user stay continuously oriented in a complex body of materials.

This is a hypertext design that can be easily understood-- if not from reading this note, from a few minutes of exploring the text at a CRT.

The design consists of text arranged and interlinked in three different sections; a special user dashboard; and a supporting system of computer and program to make the overall text object respond in the ways suggested.

This design would not be hard to implement right now on a high-performance CRT console with appropriately large memory behind the computer.

These parts are inseparable: the text and links, the dashboard, the computer program. However, the program would presumably be general enough to use with other hypertext designs as well.

The overall purpose of the design is to keep the reader oriented as to where he is and what he is doing. It is intended that the text should be under the full control of the user, that he may quickly find answers to his questions, and find interesting cross-cutting collateral material as well.

As I see it this design is best suited to discursive materials, in particular history or historically-taught survey courses. But in principle anything that can be put in a book can be put into this hypertext design. It has, however, specific editorial properties.

THE TEXT.

The text has three parts.

1) The Narrative, a unit of stretchtext (described elsewhere). At its minimum it is (say) a thousand words. When stretched it gets much larger.

2) The Lookup. This consists of definitions of words (from the dictionary), and stretchtext expositions of certain other terms, and stretchtext biographies. To reach the lookup from any place in the whole hypertext, the user has merely to point to a word, term or name, and press "jump." This definition or exposition then appears on the screen. Because these are stretchtext, if the material is not adequately clear to the reader at the first look, he may magnify it.

3) The Articles-- text pieces ranging from, say, 100 to 10,000 words in length. They may be either ordinary straight text or stretchtext. These contain "collateral material"-- insights, wisecracks, essays-- that tie together specific aspects or ideas also found in the Narrative. These articles may or may not have titles, headings, author credits, etc.

JUMPS.

The text is dotted with jump markers. A jump marker is like an asterisk: a symbol informing the reader that a branch is available at this point. To follow the jump marker,

one has merely to point at the jump marker with light pen or other device, and press the "jump" button.

There are two kinds of jump markers, which we may call N and A. N signifies an optional jump to part of the Narrative; A signifies a jump to an Article. Jump markers of both type N and A may appear anywhere in the hypertext. (The further stretchtext is magnified, the more jump markers appear.) A type N jump marker takes the reader to some place in the Narrative, where the type A takes him to the beginning (or other part) of some Article. (There probably need to be "altitude" restrictions on this that I won't go into here.)

INDICATORS AND CONTROLS.

The dashboard of the system is intended to keep the user aware of where he is, and keep him in control. To keep him aware of where he is, there are three indicator lights, representing the three parts of the text. If he goes from Articles to Narrative, or Narrative to Lookup, he may verify this immediately from the indicator lights.

The user has two simple (but sophisticated) controls. One is a throttle handle which can be moved, by degrees, in four different directions, corresponding to movement in the text: right (forward), left (back), down (magnify), up (shrink). Movements of the throttle effect corresponding movements of the text.

The other control is an Englebart-type mouse-- a box on wheels, whose movements on the table control the movements of a pointer on the screen. There are two buttons on the mouse: "jump" and "return." Various things will happen if you press "jump." If you are pointing at a word, term or name, you will get a definition, exposition or biography. If you are pointing at a jump marker, you will jump to the linked location in the Narrative or Articles.

The RETURN button undoes your jumps, popping them one at a time from a push-down list. Thus no matter how much you jump, it will be fairly easy to keep track of where you are, and get back to where you were.

This hypertext design could be easily made plain to a seven-year-old, would permit great personal option and control, and make possible an extremely rich compound of interesting material in which the user could easily stay oriented.

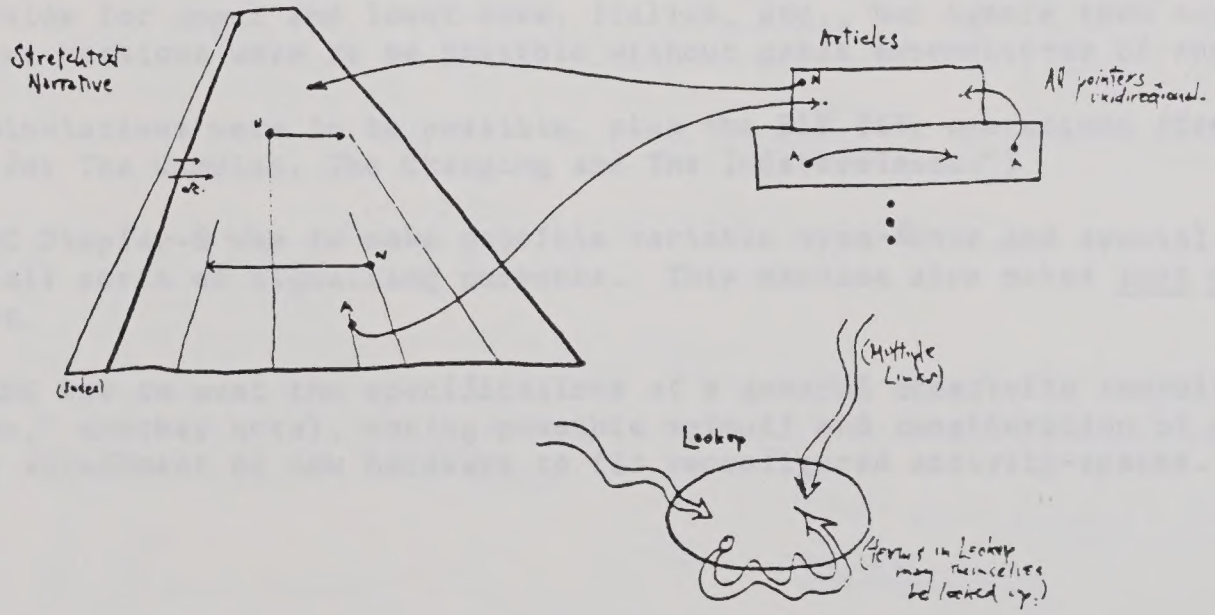

(30)

(17)

XANADU
Theodor H. Nelson

Hypertext Note 10. 29 Apr 67
Private circulation only. This will eventually be incorporated in a longer article.
Terminology is tentative.

XANADU is to text diaplay systems what SOLOMON was to big computers: an ambitious proposal, dropped by management, which may yet serve to focus thought. Indeed, it may be thought of as exemplifying a class of systems which, unless quite impractical in various technical respects, might turn out to be a prototype for our business, educational, artistic and home text display systems of the future.

XANADU was intended to do text editing; serve as a file handling, information retrieval and management information system; make possible the swift creation of activity spaces for any purpose; and serve for the viewing and creation of hypertexts, as well as lesser text objects.

The general philosophy was that the design of user front ends should be more flexible. In many systems most of the work goes into disk I/O, list processing, etc., and the screen activity design just more or less happens accidentally.

XANADU intended to reverse this, permitting the creation and maintenance of data in any structure, and the ready creation of screen formats and user activity-spaces quickly, at the console, with swift changes possible. The subtlety of screen-work more complex than, say, airline reservations makes it likely that swift change of these activity spaces--'debugging,' upgrading-- will be necessary. Personal idiosyncrasy may also dictate changes and variants in activity spaces. But activity spaces are now only modified at the machine-language level, with quick and simple reconfigurability unheard of. This was the most basic and hardest part. Virtual pushbuttons could be swiftly created, given functions, and then phased out in favor of hardware keyboards. Windows and formats could be quickly made, rearranged.

Files. Entries of any length were to have been possible with any number of connector types, and alternative sequences available. (This alternative sequence was to be related to a 'scatter write' facility, by which materials having given formal attributes could be splattered out onto the screen in different ways. (This we might call 'pseudo-sequence.') Storage was to provide for upper and lower case, italics, etc., but ignore them during search. Finally, alternative versions were to be possible without great expenditures of space.

Many text manipulations were to be possible, plus the ELF file operations (described in "A File Structure for The Complex, The Changing and The Indeterminate.")

Use of the DEC Display-8 was to make possible variable type-fonts and special characters in profusion, for all sorts of signalling purposes. This machine also makes text motions easier to implement.

Finally, XANADU was to meet the specifications of a general creativity console (see "creativity systems," another note), making possible spinoff and consideration of alternative versions and ready attachment of new hardware to fit reconfigured activity-spaces.

LECTURE HANDOUT: 'Electronic Literature and the Windowing Principle.'

Theodor H. Nelson, 28 April 1976

COMPETITIVE ECOSYSTEMS in next five or so years: CAI 'PLATO'; RCA 'Tied'; IBM satellite/ustwt-4-pcs network; Xerox 'Office of the future'; Xerox 'Dynabook'; Engelbart NLS; dreams of 'Information Retrieval,' 'Computer-Assisted Instruction,' 'Artificial Intelligence'; Video Disks; home computer with emulation games & graphics; Xanadu™ literary system:

A PROPOSED ELECTRONIC LITERARY SYSTEM & NETWORK, using computer display & transmission.

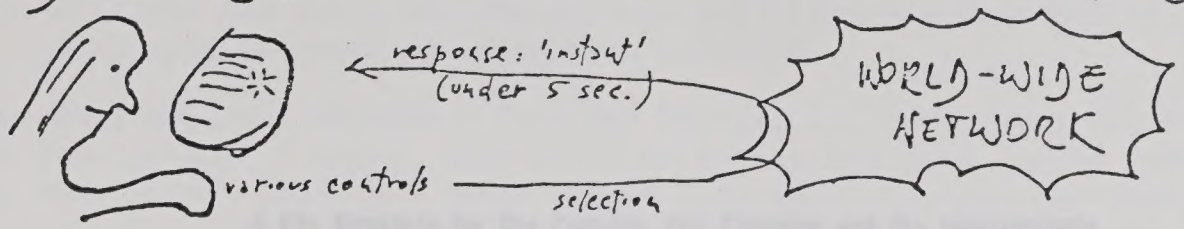

Purpose: same as reading & writing today. Fiction, scholarship, humor, directions...

Materials have three levels of accessibility: PRIVATE (others mayn't see); OPEN (not advertisable); PUBLISHED.

Every reader may instantly become an author. The author may:

- create text ('writing') } link & write ('annotation')
- link text
- excise text (intellogizing).

THE LINK MODES.

The heart of this proposal is in the exact nature of the link modes, which correspond to aspects of present-day writing. These modes also work between pictures, or between text and pictures.

1. The jump-link. Similar to footnote.

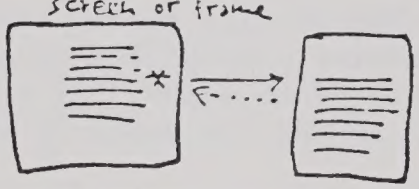

Asterisk shows user that he may jump to other material. When done, he may jump back.

2. "Collateration." (Multiple ties, between texts.) Similar to marginal notes, comparisons.

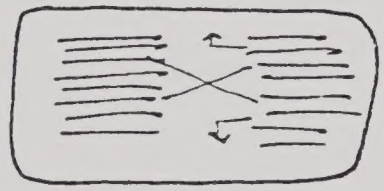

Visible lines (or broken arrows) indicate corresponding parts.

3. Quote-window. Allows jump to original source context of cited material. Similar to looking it up.

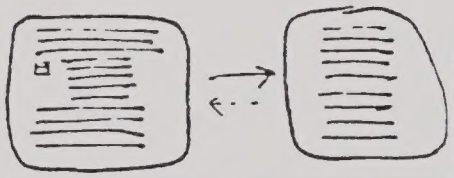

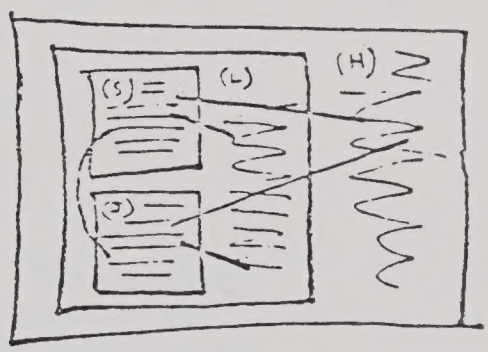

EXAMPLE OF COMPOUND WINDOWING: Sam and Joe have written documents. Louise has collaterated them with comments. Harold comments on all of those in a larger document.

All may be moved rapidly in their respective windows, with the links adjusting.

WHY THESE STRUCTURES, which I tentatively propose?

They are very powerful.
They are RELATIVELY CLEAR, compared to most alternatives.
They retain all the traditions of a pluralistic literary system — each author's work competes on its own merits.
All existing literary structures can be mapped into this, and then some, like interactive animations.
They are indefinitely extensible;
They RETAIN OBSCURENESS TO ANY DEPTH.

Each document is 'owned' — system does not mush anything together.

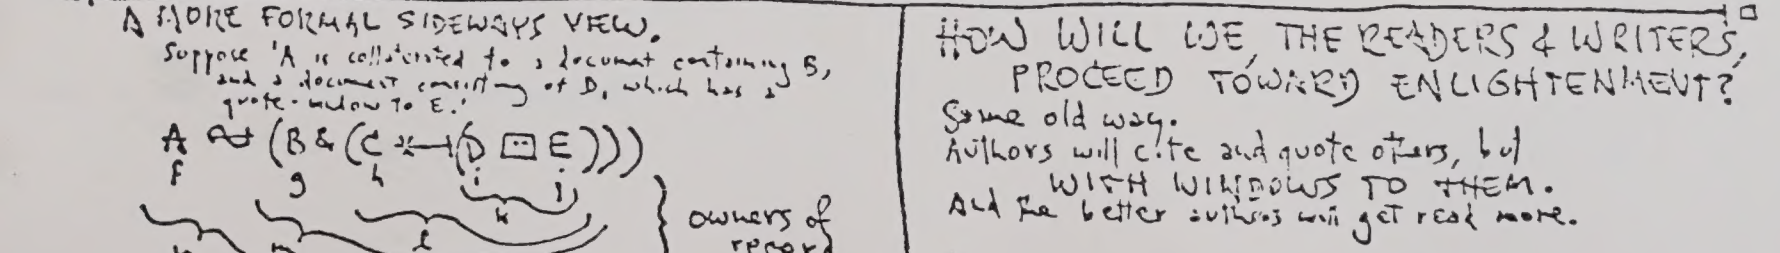

A MORE FORMAL SIDEWAYS VIEW.

Suppose 'A' is collaterated to a document containing B, and a document consisting of D, which has a quote-window to E.

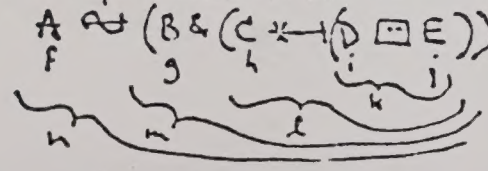

HOW WILL WE, THE READERS & WRITERS, PROCEED TOWARD ENLIGHTENMENT?

Some old way.
Authors will cite and quote others, but WITH WINDOWS TO THEM.
And the better authors will get read more.

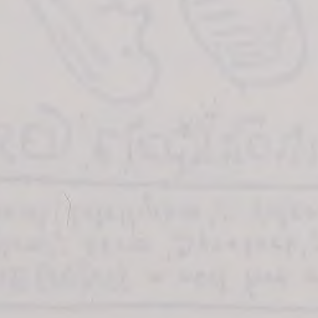

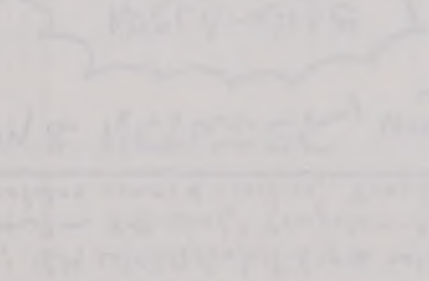

(32)

A File Structure for The Complex, The Changing and the Indeterminate

Theodor H. Nelson
Vassar College, Poughkeepsie, N.Y. 12601

Reprinted from
Association for Computing Machinery Proceedings of the 20th National Conference
with permission of ACM and the editor-publisher, Lewis Winner.
Minor corrections by the author.

© 1965 By The
Association for Computing Machinery

(33)

SUMMARY

THE KINDS OF FILE structures required if we are to use the computer for personal files and as an adjunct to creativity are wholly different in character from those customary in business and scientific data processing. They need to provide the capacity for intricate and idiosyncratic arrangements, total modifiability, undecided alternatives, and thorough internal documentation.

The original idea was to make a file for writers and scientists, much like the personal side of Bush's Memex, that would do the things such people need with the richness they would want. But there are so many possible specific functions that the mind reels. These uses and considerations become so complex that the only answer is a simple and generalized building-block structure, user-oriented and wholly general-purpose.

The resulting file structure is explained and examples of its use are given. It bears generic similarities to list-processing systems but is slower and bigger. It employs zippered lists plus certain facilities for modification and spin-off of variations. This is technically accomplished by index manipulation and text patching, but to the user it acts like a multifarious, polymorphic, many-dimensional, infinite blackboard.

The ramifications of this approach extend well beyond its original concerns, into such places as information retrieval and library science, motion pictures and the programming craft; for it is almost everywhere necessary to deal with deep structural changes in the arrangements of ideas and things.

I want to explain how some ideas developed and what they are. The original problem was to specify a computer system for personal information retrieval and documentation, able to do some rather complicated things in clear and simple ways. The investigation gathered generality, however, and has eventuated in a number of ideas. These are an information structure, a file structure, and a file language, each progressively more complicated. The information structure I call zippered lists; the file structure is the ELF, or Evolutionary List File; and the file language (proposed) is called PRIDE.

In this paper I will explain the original problem. Then I will explain why the problem is not simple, and why the solution (a file structure) must yet be very simple. The file structure suggested here is the Evolutionary List File, to be built of zippered lists. A number of uses will be suggested for such a file, in order to show the breadth of its potential usefulness. Finally, I want to explain the philosophical implications of this approach for information retrieval and data structure in a changing world.

I began this work in 1960, with no help from anybody. Its purpose was to create techniques for handling personal file systems and manuscripts in progress. These two purposes are closely related and not sharply distinct. Many writers and research professionals have files or collections of notes which are tied to manuscripts in progress. Indeed, often personal files shade into manuscripts, and the assembly of textual notes becomes the writing of text without a sharp break.

I knew from my own experiment what can be done for these purposes with card file, notebook, index tabs, edge-punching, file folders, scissors and paste, graphic boards, index-strip frames, Xerox machine and the roll-top desk. My intent was not merely to computerize these tasks but to think out (and eventually program) the dream file: the file system that would have every feature a novelist or absent-minded professor could want, holding everything he wanted in just the complicated way he wanted it held, and handling notes and manuscripts in as subtle and complex ways as he wanted them handled.

Only a few obstacles impede our using computer-based systems for these purposes. These have been high cost, little sense of need, and uncertainty about system design.

The costs are now down considerably. A small computer with mass memory and video-type display now costs $37,000; amortized over time this would cost less than a secretary, and several people could use it around the clock. A larger installation servicing an editorial office or a newspaper morgue, or a dozen scientists or scholars, could cost proportionately less and give more time to each user.

The second obstacle, sense of need, is a matter of fashion. Despite changing economies, it is fashionably believed that computers are possessed only by huge organizations to be used only for vast corporate tasks or intricate scientific calculations. As long as people think that, machines will be brutes and not friends, bureaucrats and not helpmeets. But since (as I will indicate) computers could do the dirty work of personal file and text handling, and do it with richness and subtlety beyond anything we know, there ought to be a sense of need. Unfortunately, there are no ascertainable statistics on the amount of time we waste fussing among papers and mislaying things. Surely half the time spent in writing is spent physically rearranging words and paper and trying to find things already written; if 95% of this time could be saved, it would only take half as long to write something.

The third obstacle, design, is the only substantive one, the one to which this paper speaks.

Let me speak first of the automatic personal filing system. This idea is by no means new. To go back only as far as 1945, Vannevar Bush, in his famous article "As We May Think"[1], described a system of this type. Bush's paper is better remembered for its predictions in the field of information retrieval, as he foresaw the spread and power of automatic document handling and the many new indexing techniques it would necessitate. But note his predictions for personal filing:

> Consider a future device for individual use, which is a sort of mechanized private file and library. It needs a name, and, to coin one at random, "memex" will do. A memex is a device in which an individual stores all his books, records, and communications, and which is

mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.

It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works. On the top are slanting translucent screens, on which material can be projected for convenient reading. There is a keyboard, and sets of buttons and levers. Otherwise it looks like an ordinary desk.

....

A special button transfers him immediately to the first page of the index. Any given book of his library [and presumably other textual material, such as notes] can thus be called up and consulted with far greater facility than if it were taken from a shelf. As he has several projection positions, he can leave one item in position while he calls up another. He can add marginal notes and comments, ..." (1, 106-7)

Understanding that such a machine required new kinds of filing arrangements, Bush stressed his file's ability to store related materials in associative trails, lists or chains of documents joined together.

When the user is building a trail, he names it, inserts the name in his code book, and taps it out on his keyboard. Before him are the two items to be joined, projected onto adjacent viewing positions. At the bottom of each there are a number of blank code spaces, and a pointer is set to indicate one of these on each item. The user taps a single key, and the items are permanently joined.

....

Thereafter, at any time, when one of these items is in view, the other can be instantly recalled merely by tapping a button below the corresponding code space. Moreover, when numerous items have been thus joined together to form a trail, they can be reviewed in turn, rapidly or slowly, by deflecting a lever like that used for turning the pages of a book. It is exactly as though the physical items had been gathered together from widely separated sources and bound together to form a new book. It is more than this, for any item can be joined into numerous trails.

....

Thus he goes, building a trail of many items. Occasionally he inserts a comment of his own, either linking it into the main trail or joining it by a side trail to a particular item. (1, 107)

Two decades later, this machine is still unavailable*.

The hardware is ready. Standard computers can handle huge bodies of written information, storing them on magnetic recording media and displaying their contents on CRT consoles, which far outshine desktop projectors. But no programs, no file software are standing ready to do the intricate filing job (keeping track of associative trails and other structures) that the active scientist or thinker wants and needs. While Wallace[3] reports that the System Development Corporation has found it worthwhile to give its employees certain limited computer facilities for their own filing systems, this is a bare beginning.

Let us consider the other desideratum, manuscript handling. The remarks that follow are intended to apply to all forms of writing, including fiction, philosophy, sermons, news and technical writing.

* The Busn Rapid Selector, which he designed (2), is a powerful microfilm instrument, but it is not suited to idiosyncratic personal uses, nor to evolutionary modification, as described hereunder.

The problems of writing are little understood, even by writers. Systems analysis in this area is scanty; as elsewhere, the best doers may not understand what they do. Although there is considerable anecdote and lore about the different physical manuscript and file techniques of different authors, literary tradition demerits any concern with technical systems as detracting from "creativity." (Conversely, technical people do not always appreciate the difficulty of organizing text, since in technical writing much of the organization and phraseology is given, or appears to be.) But in the computer sciences we are profoundly aware of the importance of systems details, and of the variety of consequences for both quality and quantity of work that result from different systems. Yet to design and evaluate systems for writing, we need to know what the process of writing is.

There are three false or inadequate theories of how writing is properly done. The first is that writing is a matter of inspiration. While inspiration is useful, it is rarely enough in itself. "Writing is 10% inspiration, 90% perspiration," is a common saying. But this leads us to the second false theory, that "writing consists of applying the seat of the pants to the seat of the chair." Insofar as sitting facilitates work, this view seems reasonable, but it also suggests that what is done while sitting is a matter of comparative indifference; probably not.

The third false theory is that all you really need is a good outline, created on prior consideration, and that if the outline is correctly followed the required text will be produced. For most good writers this theory is quite wrong. Rarely does the original outline predict well what headings and sequence will create the effects desired: the balance of emphasis, sequence of interrelating points, texture of insight, rhythm, etc. We may better call the outlining process inductive: certain interrelations appear to the author in the material itself, some at the outset and some as he works. He can only decide which to emphasize, which to use as unifying ideas and principles, and which to slight or delete, by trying. Outlines in general are spurious, made up after the fact by examining the segmentation of a finished work. If a finished work clearly follows an outline, that outline probably has been hammered out of many inspirations, comparisons and tests**.

Between the inspirations, then, and during the sitting, the task of writing is one of rearrangement and reprocessing, and the real outline develops slowly. The original crude or fragmentary texts created at the outset generally undergo many revision processes before they are finished. Intellectually they are pondered, juxtaposed, compared, adapted, transposed, and judged; mechanically they are copied, overwritten with revision markings, rearranged and copied again. This cycle may be repeated many times. The whole grows by trial and error in the processes of arrangement, comparison and retrenchment. By examining and mentally noting many different versions, some whole but most fragmentary, the intertwining and organizing of the final written work gradually takes place***.

Certain things have been done in the area of computer manuscript handling. IBM recently announced its "Administrative Terminal System"[5,6,7,8] which permits the storage of unfinished sections of text in computer memory, permits various modifications by the user, and types up the final draft with page numbers, right justification and headers.

While this is a good thing, its function for manuscripts is cosmetic rather than organizing. Such a system can be used only with textual sections which are already well organized, the visible part of the iceberg. The major and strenuous part of such writing must already have been done.

** I understand that this account is reasonably correct for such writers as Tolstoy, Winston Churchill and Katherine Anne Porter. Those who can stick to a prior outline faithfully, like James Fenimore Cooper, tend to be either hacks or prodigies, and don't need this system.

*** For a poignant, mordant portrayal of the writer's struggle, the reader is directed to Gorey's The Unstrung Harp; or, Mr Earbrass Writes a Novel.

If a writer is really to be helped by an automated system, it ought to do more than retype and transpose: it should stand by him during the early periods of muddled confusion, when his ideas are scraps, fragments, phrases, and contradictory overall designs. And it must help him through to the final draft with every feasible mechanical aid-- making the fragments easy to find, and making easier the tentative sequencing and juxtaposing and comparing.

It was for these two purposes, taken together-- personal filing and manuscript assembly-- that the following specifications were drawn up.

Here were the preliminary specifications of the system: It would provide an up-to-date index of its own contents (supplanting the "code book" suggested by Bush). It would accept large and growing bodies of text and commentary, listed in such complex forms as the user might stipulate. No hierarchical file relations were to be built in; the system would hold any shape imposed on it. It would file texts in any form and arrangement desired-- combining, at will, the functions of the card file, loose-leaf notebook, and so on. It would file under an unlimited number of categories. It would provide for filing in Bush trails. Besides the file entries themselves, it would hold commentaries and explanations connected with them. These annotations would help the writer or scholar keep track of his previous ideas, reactions and plans, often confusingly forgotten.

In addition to these static facilities, the system would have various provisions for change. The user must be able to change both the contents of his file and the way they are arranged. Facilities would be available for the revising and rewording of text. Moreover, changes in the arrangements of the file's component parts should be possible, including changes in sequence, labelling, indexing and comments.

It was also intended that the system would allow index manipulations which we may call dynamic outlining (or dynamic indexing). Dynamic outlining uses the change in one text sequence to guide an automatic change in another text sequence. That is, changing an outline (or an index) changes the sequence of the main text which is linked with it. This would permit a writer to create new drafts with a relatively small amount of effort, not counting rewordings.

However, because it is necessary to examine changes and new arrangements before deciding to use or keep them, the system must not commit the user to a new version until he is ready. Indeed, the system would have to provide spin-off facilities, allowing a draft of a work to be preserved while its successor was created. Consequently the system must be able to hold several-- in fact, many-- different versions of the same sets of materials. Moreover, these alternate versions would remain indexed to one another, so that however he might have changed their sequences, the user could compare their equivalent parts.

Three particular features, then, would be specially adapted to useful change. The system would be able to sustain changes in the bulk and block arrangements of its contents. It would permit dynamic outlining. And it would permit the spin-off of many different drafts, either successors or variants, all to remain within the file for comparison or use as long as needed. These features we may call evolutionary.

The last specification, of course, one that emerged from all the others, was that it should not be complicated.

These were the original desiderata. It was not expected at first that a system for this purpose would have wider scope of application; these jobs seemed to be quite enough. As work continued, however, the structure began to look more simple, powerful and general, and a variety of new possible uses appeared. It became apparent that the system might be suited to many unplanned applications involving multiple categories, text summaries or other parallel documents, complex data structures requiring human attention, and files whose relations would be in continuing change.

Note that in the discussion that follows we will pretend we can simply see into the machine, and not worry for the present about how we can actually see, understand and manipulate these files. These are problems of housekeeping, I/O and display, for which many solutions are possible.

Elements of the ELF

What was required we may call an evolutionary file structure: a file structure that can be shaped into various forms, changed from one arrangement to another in accordance with the user's changing need. It was apparent also that some type of list structure was necessary. Making the file out of lists would allow different categories of personal notes, separate drafts, outlines and master indices all to be handled as lists of some sort; their segments could then be manipulated through automatic handling of index numbers. The resulting file structure I will accordingly call the Evolutionary List File, or ELF, since it is an evolutionary file structure constructed with lists. The system proposed here is not the only ELF possible. It is built upon a specific technique of attaching lists together which has a natural resistance to becoming confused and messy.

As computer-based systems grow in capability and diversity of uses, they tend to become more and more cluttered with niggling complications, hidden passageways, and lurking, detailed interlocks, restrictions, specializations, provisos. These should be forsworn, if possible, in the system under discussion, so that it might be attractive to laymen (including artists and writers) who feel unkindly disposed toward computers. It should readily adapt to their own styles of handling things, imposing few conventions or methods of use. How could this imposition be avoided? And among so many interesting and possible system functions and file relations, how may the users know what connections to make, how may they understand what they are doing, and how may they avoid muddling and losing the things they are working with?

The answer, I think you see, is to choose a very simple structure that can be used and compounded in many different ways. The basic arrangement chosen for these purposes is an information structure I will refer to as zippered lists. (We might call it permutation-invariant one-for-one inter-list entry-linking, but that is not necessary.)

There are only three kinds of things in the zippered-list ELF, with no predetermined relations among them-- no hierarchies, machine-based features or trick exceptions. The system is user-oriented and open-faced, and its clear and simple rules may be adapted to all purposes.

The ELF has three elements: entries, lists and links. An entry is a discrete unit of information designated by the user. It can be a piece of text (long or short), a string of symbols, a picture or a control designation for physical objects or operations.

A list is an ordered set of entries designated by the user. A given entry may be in any number of lists.

A link is a connector, designated by the user, between two particular entries which are in different lists (Figure 1). An entry in one list may be linked to only one entry in another list.

On the left we see two zippered lists. Between the entries of list A and those of B are dashed lines, representing the links between the two lists. On the right is the table of links as it might look to a machine. The machine can read this table from right to left or left to right, finding entries in B that correspond to given entries in A, or vice versa. A change in the sequence of either list, or additions to either list, will not change the links that stand between them. Changes in the link structure will occur only if the user specifically changes the links, or if he destroys entries which are linked to others.

To be technical, then, two lists are zippered if there are any pairwise links between their respective elements, each element is in no more than one link pair, and these links are unaffected by permutation of the lists, remaining affixed to the same pairs of elements. It is not required that the two lists be of the same length, or, even if they are, that all entries have a link to the other list.

The ELF's File Operations

Zippered lists are an information structure; the Evolutionary List File is a file structure. The ELF described in this paper holds its contents exclusively as zippered (or unzippered) lists. But the file structure must also include a set of operations by which it may be modified. These file operations exist for creating, adjusting or removing the entry, list and link, and for manipulating the sequence relation. An ELF is actually any machine which will, on command, carry out the basic operations on entry, list, link and sequence.

Entries. The user may create new entries at any time, putting anything in them that he thinks appropriate. Entries may be combined or divided (unless indivisible, like objects, commands, etc.) Entries may be put in any list, and the same entry may be put in different lists. The user may direct that entries of one list be automatically copied onto another list, without affecting the original list.

Lists. The user may create lists and assign entries to them. He may at will make new copies of lists. He may rearrange the sequence of a list, or copy the list and change the sequence of that copy. Lists may be combined; lists may be cut into sublists.

Links. The user may create links between entries that are in different lists. Any number of legal links may be created, although the upper limit of links between any two lists is determined by the 1-for-1 rule. When an entry or a list is copied into a list, links will remain between parent and daughter entries. Moreover, after a list-copying operation, the daughter list will have the same links to all other lists as does the parent list.

Sequences. The user may put a list in any sequence he wishes. (A copied list will maintain the original sequence until modified.) Sequences may be transferred between lists via the links: if the sequence of A is transferred to B, each entry of A linked to an entry in B takes the sequential position of its linked

entry in B.

No definite meaning is assigned to these entities or operations by the system; the user is free to let them mean anything he likes. A list may be a category, trail, index, dialogue, catalogue or poem, and lists may be assembled into larger structures. The ELF may be thought of as a place; not a machine, but a piece of stationery or office equipment with many little locations which may be rearranged with regard to one another****.

Note that zippered lists generate only one of various possible Evolutionary List Files. Indeed, the description of the file structure given here is in some ways restrictive: the ELF could take a number of other, closely similar forms and still be much the same thing. For example, it would be possible to allow sub-entries and superentries into the file, to behave and link up like normal entries, even though they contained or were contained in other entries. But the equivalent can be done with the current system. Another possibility would be to allow links other than 1-for-1; these could be modal, the different link-modes having different meanings to the user. Or we might make it an evolutionary network file, allowing any two entries to be connected.. Or, besides such general changes in the rules, plausible changes and accessory functions for any purposes could be introduced outside the given file structure, even including modifications and widgets to do some of the same things "more easily."

But to the user such complication might render the system far less handy or perspicuous. The ELF, with its associated techniques as described above, is simple and unified. Many tasks can be handled within the file structure. This means it can be of particular benefit to people who want to learn without complications and use it in ways they understand. For psychological, rather than technical reasons, the system should be lucid and simple. I believe that this ELF best meets these requirements.

Technical Aspects

Since the ELF description above bears some resemblance to the list languages, such as IPL, SLIP, etc., a distinction should be drawn. These list languages[9] are particularly suited to processing data, fast and iteratively, whose elements are manipulable in Newell-Shaw-Simon lists. Essentially they may be thought of as organizations of memory which facilitate sequential operations on unpredictably branching or hierarchical data. These data may change far too quickly for human intervention. Evolutionary file structures, and the ELF in particular, are designed to be changed piecemeal by a human individual. While it might be convenient to program an ELF in one of these languages, the low speed at which user file commands need to be executed makes such high-powered implementation unnecessary; the main problem is to keep track of the file's arrangements, not to perform computation on its contents. Although work has been done to accommodate the list-language approach to larger chunks of material than usual[10], the things people will want to put into an ELF will typically be too big for core memory.

The ELF does in fact share some of the problems of the list languages: not available-storage accounting or garbage collection (concerns associated with organization of fast memory for processing, which may be avoided at slower speeds), but the problems of checkout for disposal (what other lists is an entry on?) and list naming. The former problem is rather straightforwardly solved[11, p. 164]; the latter is complicated in ways we cannot go into here.

**** An ELF might even be constructed out of cards, blocks, sticks and strings, using techniques of puppetry, but this would not be a convenient object.

The ELF appears to be closest, topologically and in other organizing features, to the Multilist system described by Prywes and Gray[12]. Like that system, it permits putting entries in many different lists at once. However, in current intent[13] that system is firmly hierarchical, and thus somewhat removed from the ELF's scope of application. Another closely related system is the Integrated Data Store of Bachman[14,15,16,17,18]; this is intended as a hardware-software system for disc I/O and storage arrangement, but in its details it seems the ELF's close relative. Each of these systems has a connection logic that might be feasible as a basis for an ELF different from this one. Or, either might prove a convenient programming base for the implementation of this file structure.

Another obvious technical question must be considered. How can the ELF allow "unlimited" copies of entries and lists? By patching techniques, of course. Variant entries and lists can take virtually no space, being modification data plus pointers to the original. When a modified version of a list or entry is created, the machine patches the original with the changes necessary to make the modified version. (Figure 2)

USES

In the discussion that follows, we will examine various possible applications of zippered lists and the ELF, and postpone discussing the file language they require. Finally we will return to this problem, and describe the file language PRIDE whose additional features are needed to adapt the ELF for the uses originally discussed.

By assigning entries to lists, the ELF may be used as a glorified card file, with separate lists used for categories, trails, etc. This permits extensive cross-indexing by the assignment of one entry to different lists. It permits subsets and sub-sequences for any use to be held apart and examined without disturbing the lists from which they have been drawn, by copying them onto other, new lists. The ELF permits the filing of historical trails or associative (Bush) trails through documents, business correspondence, belles-lettres, case law, treaties, scholarly fields and history, and the mixture of trail with categorical filing.

These are the simple uses; the compound uses are much more interesting. But since we cannot intuitively fit every possible conceptual relationship into zippered lists, imaginative use is necessary. Remember that there is no <u>correct</u> way to use the system. Given its structure, the user may figure out any method useful to him. A number of different arrangements can be constructed in the ELF, using only the basic elements of entry, list and link. Zippered lists may be assembled into rectangular arrays, lattices and more intricate configurations. These assemblies of lists may be assigned meaning <u>in combination</u> by the user, and the system will permit them to be stored, displayed, taken apart for examination, and corrected, updated, or modified.

By using such combining arrangements on lists composed of text, the file can be self-documenting, with all labelling and documentation kept integrally within the file structure. It is thus possible to incorporate, in a body of information filed in the ELF, various levels of index, summary, explanation and commentary. Many useful ways of listing and linking such documentation are possible. In Figure 3 we see some of the ways that documentary lists may be linked together. The lists shown are outline, suboutline, draft, subdraft, summary, commentary and source list. These are not all the possible types of documentary lists; for example, "footnotes" are omitted. The ELF will permit any number of these documentary lists;

observe that they can be built on one another, and indefinitely compounded. The system will have no trouble accepting a commentary on a commentary on a subdraft of an outline for a variant list of source materials.

Figure 3 shows also how two lists may contain some of the same entries. The dashed line represents linkage between entries, the solid line shows that both lists contain the same entry. This may be useful for creating alternate versions, or, as in this example, the lists containing the same entry may have different purposes. Here, for instance, an entry in the summary is also to be found in the main draft.

This self-documentation feature permits any string of text in the ELF, long or short, to be annotated or footnoted for scholarly or other purposes. Such marginalia can be temporary or permanent, for the private memoranda of an individual or for communication among different persons using the file.

In a like manner, the ELF is capable of storing many texts in parallel, if they are equivalent or linked in some way. For example, instruction manuals for different models of the same machine may be kept in the file as linked lists, and referred to when machines are to be compared, used or fixed. This is of special use to repairmen, project managers and technical writers.

Moreover, the ELF's cross-sequencing feature -- the fact that links ignore permutations-- permits the collation of very different cognate textual materials for comparison and understanding. In law, this would help in comparing statutes (or whole legal systems); in literature, variorum editions and parodies. Thus such bodies as the Interpreter's Bible and a Total Shakespeare (incorporating Folios, bowdlerizations, satires and all critical commentary) could be assembled for study.

Let me try to illustrate the possible comprehensiveness and versatility of this file structure as applied to texts. Figure 4 shows the different arrangements that might be used by one man-- in this case an historian writing a book-- to assemble and integrate his intellectual and professional concerns. Although it is impossible to show the links between all the separate entries of these lists-- the entries are not themselves discernible in this drawing-- it is possible to note the kinds of links between lists. A thin line between lists shows that some links exist; a solid line indicates that some entries of both lists are the same.

Perhaps this looks complicated. In fact, each of the connectors shows an indexing of one body of information to another; this user may query his file in any direction along these links, and look up the parts of one list which are related to parts of another. Therefore the lines mean knowledge and order. Note that in such uses it is the man's job to draw the connections, not the machine's. The machine is a repository and not a judge.

The ELF may be an aid to the mind in creative tasks, allowing the user to compare arrangements and alternatives with some prior ideal. This is helpful in planning nonlinear assemblages (museum exhibits, casting for a play,) or linear constructions of any kind. Such linear constructions include not only written texts; they can be any complicated sequences of things, such as motion pictures (in the editing stage) and computer programs.

Indeed, computer programming with an on-line display and the ELF would have a number of advantages. Instructions might be interleaved indefinitely without

resorting to tiny writing. Moreover, the programmer could keep up work on several variant approaches and versions at the same time, and easily document their overall features, their relations to one another and their corresponding parts. Adding a load-and-go compiler would create a self-documenting programming scratchpad.

The natural shape of information, too, may call for the ELF. For instance, sections of information often arrange themselves naturally in a lattice structure, whose strands need to be separately examined, pondered or tested. Such lattices include PERT networks, programmed instruction sequences, history books and genealogical records. (The ELF can handle genealogical source documentation and its original text as well.) Indeed, any informational networks that require storage, handling and consideration will fit the ELF; a feature that could have applications in plant layout, social psychology, contingency planning, circuit design and itineraries.

The ELF may, through its mutability, its expansibility, and its self-documentation features, aid in the integration, understanding and channeling of ideas and problems that will not yield to ordinary analysis or customary reductions; for instance, the contingencies of planning, which are only partially Boolean. Often the reason for a so-called Grand Strategy in a setting is that we cannot keep track of the interrelations of particular contingencies. The ELF could help us understand the interrelations of possibilities, consequences, and strategic options. In a logically similar case, evaluating espionage, it might help trace consistencies and contradictions among reports from different spies.

The use of an ELF as the basis for a management information system is not inconceivable. Its evolutionary capability would provide a smooth transition from the prior systems, phasing out old paperwork forms and information channels piecemeal. Beginning with conventional accounting arrays and information flow, and moving through discrete evolutionary steps, the ELF might help restructure an entire corporate system. Numerical subroutining could permit the system to encompass all bookkeeping. The addresses of all transaction papers, zippered to lists of their dates and contents, would aid in controlling shipments, inventory and cash. The ELF's cross-sequencing feature could be put to concrete uses, helping to rearrange warehouses (and the company library) by directing the printout of new labels to guide physical rearrangement. Inventories, property numbers and patents could be so catalogued and recatalogued in the ELF. Legal documents, correspondence, company facts and history could be indexed or filed in historical and category trails. And upper management could add private annotations to the public statements, reports and research of both the organization and its competitors, with amendments, qualifications, and inside dope.

PRIDE

While the ELF as described is expected to be general and useful, the original purposes described at the beginning of this paper call for certain further provisions. Now I would like to describe a desirable file and information handling language that will meet these needs, called the PRIDE (Personalized Retrieval, Indexing, and Documentation Evolutionary) System. Its purpose is to facilitate the use of an ELF. The system described is not yet implemented, nor even fully specified, but let us speak as though it is.

PRIDE includes the ELF operations. However, for safety and convenience nearly every operation has an inverse. The user must be permitted, given a list of what he has done recently, to undo it. It follows that "destroy" instructions must

fail safe; if given accidentally, they are to be revocable. For safety's sake, it should take several steps to throw a thing away completely. An important option would permit the user to retrace chronologically everything he does on the system.

Most of PRIDE's applications will involve text handling, either as a primary purpose or in the documentation of some other task. Hence a number of features exist for convenient text usage. Text handling commands (for modifying entries) include the equivalents of standard proofreader's marks for insertion, deletion and switching of sections.

Also for text usage and user comfort, there are certain system non-restrictions. There is no practical restriction on the length of an input entry, and it need follow only the most trivial format conventions. In addition, the machine will interrupt any other PRIDE function to receive input text (inspiration mode). It is necessary that entries of unspecified length be acceptable to the system without fuss or warning. PRIDE does not stipulate fixed record lengths, either for input or storage; any such restrictions would have a psychologically cramping effect. There is no reason the system cannot appear to the user to have no fixed or standard unit lengths; the machine's operating units and sections should not concern him.

Ideally, neither the length of entries, the number of lists, or any other parameter of a file is restricted by anything but the absolute size of all memory. This is a difficult requirement for the programmer. Routinely, however, the system should be able to accept entries thousands of characters long, accept hundreds of entries to a list, and accept hundreds of lists in the file. Otherwise, extraneous consideration by the user of whether there's room to add material or try out an offshoot begins to interfere with the system's use.

Although I have avoided discussing the means by which the user sees his file, PRIDE must, of course, have functions and commands for this purpose. For a CRT these include quick lookup schemes[19], preferably with moving menus and means of readily changing the hierarchy of lookup structure; as well as visual cuing and mnemonic formats, including cursor maneuvers, overlays and animated wipes and other transitions. But such glamorous features do not reduce the challenge or worth of working through a line printer, or seeking to make the system useful under a batch-processing monitor.

Many instructions aside from those already mentioned will be needed by the user; particular applications will require such operations as text lookup and integer arithmetic. And surely all the uses of the system have not been anticipated. Hence a subroutining facility is to be available, reaching to assembly language or opening into the machine's other languages. This could be used for processing the file's contents (e.g., numbers or character strings), or for creating more convenient combined operations out of the different operations dealing with file structure, input-output and text.

PRIDE is one possible way to make an ELF, or any evolutionary file structure, useful. PRIDE would be a foreground, free-standing language with the primary mission of handling files and manuscripts, as discussed at the beginning, and secondary applications in ordering and documenting other kinds of complex information. Its major use would presumably be in connection with time-shared display and information systems. But such a language is only one suggestion. Actually, there is not much reason that the ELF could not be made a standard file structure for all purposes; unused capabilities would not intrude, but would still be there if unex-

pectedly wanted. ELF systems could be built into the file capabilities of general utility software. The actual computation involved is relatively trivial, and the ELF could easily be incorporated into I/O routines or data channel languages. Even small-scale hardware implementations are not unthinkable; a control box between a typewriter and a tape recorder, for instance.

All these applications depend, of course, on the system's being actually useful, which is an empirical question. A number of possible applications have been mentioned. But, except as a crutch to man's fallible mind, is there any reason to suppose that the system has any general applicability in principle?

Philosophy

As "philosophy" I want to speak of two major things. First, complex file structures (like the ELF) make possible the creation of complex and significant new media, the hypertext and hyperfilm. Second, evolutionary file structures (like the ELF) make it possible to keep track of things that have been changing, without our awareness, all along. These include the major categories of human thought, which will go on changing.

Systems of paper have grave limitations for either organizing or presenting ideas. A book is never perfectly suited to the reader; one reader is bored, another confused by the same pages. No system of paper-- book or programmed text-- can adapt very far to the interests or needs of a particular reader or student.

However, with the computer-driven display and mass memory, it has become possible to create a new, readable medium, for education and enjoyment, that will let the reader find his level, suit his taste, and find the parts that take on special meaning for him, as instruction or entertainment.

Let me introduce the word "hypertext"***** to mean a body of written or pictorial material interconnected in such a complex way that it could not conveniently be presented or represented on paper. It may contain summaries, or maps of its contents and their interrelations; it may contain annotations, additions and footnotes from scholars who have examined it. Let me suggest that such an object and system, properly designed and administered, could have great potential for education, increasing the student's range of choices, his sense of freedom, his motivation, and his intellectual grasp******. Such a system could grow indefinitely, gradually including more and more of the world's written knowledge. However, its internal file structure would have to be built to accept growth, change and complex informational arrangements. The ELF is such a file structure.

Films, sound recordings, and video recordings are also linear strings, basically for mechanical reasons. But these, too, can now be arranged as non-linear systems-- for instance, lattices-- for editing purposes, or for display with different emphasis. (This would naturally require computer control, using the ELF or a related system, and various cartridge or re-recording devices.) The hyperfilm-- a browsable or vari-sequenced movie-- is only one of the possible hypermedia that require our attention.

So much for what we can create afresh with this structure. What about the things that have already been around awhile?

The physical universe is not all that decays. So do abstractions and categories. Human ideas, science, scholarship and language are constantly collapsing

***** The sense of "hyper-" used here connotes extension and generality; cf. "hyperspace." The criterion for this prefix is the inability of these objects to be comprised sensibly into linear media, like the text string, or even media of somewhat higher complexity. The ELF is a hyperfile.

****** I will discuss this idea at length elsewhere.

and unfolding. Any field, and the corpus of all fields, is a bundle of relationships subject to all kinds of twists, inversions, involutions and rearrangement: these changes are frequent but unpredictable. Recall that computers, once a branch of mathematics, are now their own field (but the development of fluid logic indicates a possible merger with the art of wind instruments). Social relations, psycholinguistics and psychonomics are new fields, even though they rest on no special discoveries; political economy, natural history and social ethics are gone. Within a given area, too, the subheadings of importance are in constant flux. In the social sciences, for instance, the topic headings of the nineteen-thirties now sound quaint.

While the disappearance and up-ending of categories and subjects may be erratic, it never stops; and the meaning of this for information retrieval should be clear. Last week's categories, perhaps last night's field, may be gone today. To the extent that information retrieval is concerned with seeking true or ideal or permanent codes and categories-- and even the most sophisticated "role indicator" syntaxes are a form of this endeavor-- to this extent, information retrieval seems to me to be fundamentally mistaken. The categories are chimerical (or temporal) and our categorization systems must evolve as they do. Information systems must have built in the capacity to accept the new categorization systems as they evolve from, or outside, the framework of the old. Not just the new material, but the capacity for new arrangements and indefinite rearrangements of the old, must be possible. In this light, the ELF, indefinitely revisible and unperturbed by changes in overall structural relations, offers some promise.

There is, then, a general rationale. I believe that such a system as the ELF actually ties in better than anything previously used with the actual processes by which thought is progressively organized, whether into stories or hypertext or library categories. Thus it may help integrate, for human understanding, bodies of material so diversely connected that they could not be untangled by the unaided mind. For both logistic and psychological reasons it should be an important adjunct to imaginative, integrating and creative enterprises. It is useful where relationships are unclear; where contingencies and tasks are undefined and unpredictable; where the structures or final outcome it must represent are not yet fully known; where we do not know the file's ultimate arrangement; where we do not know what parts of the file are most important; or where things are in permanent and unpredictable flux. Perhaps this includes more places than we think. And perhaps here, as in biology, the only ultimate structure is change itself.

CONCLUSION

This paper has proposed a different kind of structure for handling information.

Essentially it is a file with certain storage provisions which, combined, permit the file's contents to be arranged any-which-way, and in any number of ways at once. A set of manipulation functions permits making changes or keeping track of developments. The file is capable of maintaining many different arrangements at the same time, many of which may be dormant. This makes ordinary measures of efficiency inappropriate; as with high fidelity music systems, enrichment is derived from the lavish use of surplus capacity.

The key ideas of the system are the inter-linking of different lists, regardless of sequence or additions; the re-configurable character of a list complex into any humanly conceivable forms; and the ability to make copies of a whole list, or

list complex-- in proliferation, at will-- to record its sequence, contents or arrangement at a given moment. The Evolutionary List File is a member of the class of evolutionary file structures; and its particular advantages are thought to be psychological, not technical. Despite this file's adaptability to complex purposes, it has the advantage of being conceptually very simple. Its structure is complete, closed, and unified as a concept. This is its psychological virtue. Its use can be easily taught to people who do not understand computers. We can use it to try out combinations that interest us, to make alternatives clear in their details and relationships, to keep track of developments as they occur, to sketch things we know, like or currently require; and it will stand by for modifications. It can be extended for all sorts of purposes, and implemented or incorporated in any programming language.

There are probably various possible file structures that will be useful in aiding creative thought. This one operates, as it were, on lists that hook together sideways, and their copies. There may be many more.

REFERENCES

1. Bush, Vannevar, "As We May Think." The Atlantic Monthly, 176:1 (July, 1945), 101-108.
2. Hirsch, Phil, "The Bush Rapid Selector." Datamation, June 1965, 56-7.
3. Wallace, Everett L., "Experience with EDP Support of Individuals' File Maintenance." Parameters of Information Science: Proceedings of the American Documentation Institute, v.I (American Documentation Institute, 1964), 259-261.
4. Gorey, Edward, The Unstrung Harp; or, Mr Earbrass Writes a Novel. Duell, Sloan and Pearce, N.Y., 1953.
5. International Business Machines Data Processing Division, "The IBM Administrative Terminal System." Company brochure 520-1146.
6. International Business Machines Technical Publications Department, "1440-/1460 Administrative Terminal System Application Description." IBM, White Plains, New York.
7. Timberlake, W.D., "Administrative Terminal System" (abstract.) STWP Proceedings, May 1965, no page number. New York: Society of Technical Writers and Publishers, 1965.
8. Farrell, Austin C., "Evolution of Automated Writing." STWP Proceedings, May 1965, no page numbers. New York: Society of Technical Writers and Publishers, 1965.
9. Bobrow, Daniel G. and Bertram Raphael, "A Comparison of List-Processing Languages." Comm. ACM 7:4 (April, 1964), 231-240.
10. Comfort, W.T., "Multiword List Items." Comm. ACM 7:6 (June, 1964), 357-362.
11. Weizenbaum, J., "Knotted List Structures." Comm. ACM 5:3 (March, 1962), 161-165.
12. Prywes, N.S. and H.J. Gray, "The Multilist System for Real Time Storage and Retrieval." Proceedings of IFIP Conference, 1962, 112-116.
13. Prywes, N.S., "Interim Technical Report: The Organization of Files for Command and Control." Moore School of Engineering, March, 1964.
14. Bachman, C.W. and S.B. Williams, "The Integrated Data Store-- A General Purpose Programming System for Random Access Memories." AFIPS Conference Proceedings, v. 26, 411-422. Spartan Books, 1964.
15. General Electric Computer Department, "I-D-S." Company brochure CPB-425.
16. General Electric Computer Department, "Introduction to Integrated Data Store." Company brochure CPB-1040, 1965.
17. Bachman, Charles W., "Software for Random Access Processing." Datamation, April 1965.
18. General Electric Computer Department, "Integrated Data Store: New General Electric Technique for Organizing Business Data." Company publication, January 1965.
19. Corbin, Harold S., and George J. Stock, "On-Line Querying via a Display Console." Fourth National Symposium on Information Display: Technical Session Proceedings, 127-134. Society for Information Display, Washington, 1964.

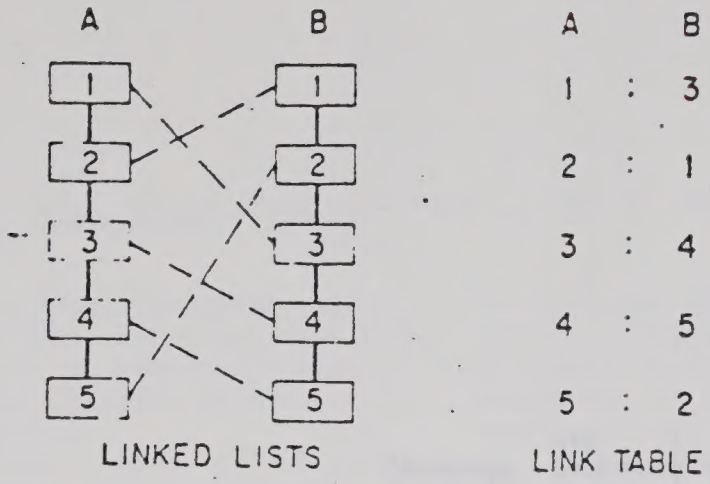

FIGURE 1—Zippered lists: 1-for-1 links between entries are invariant under list permutation.

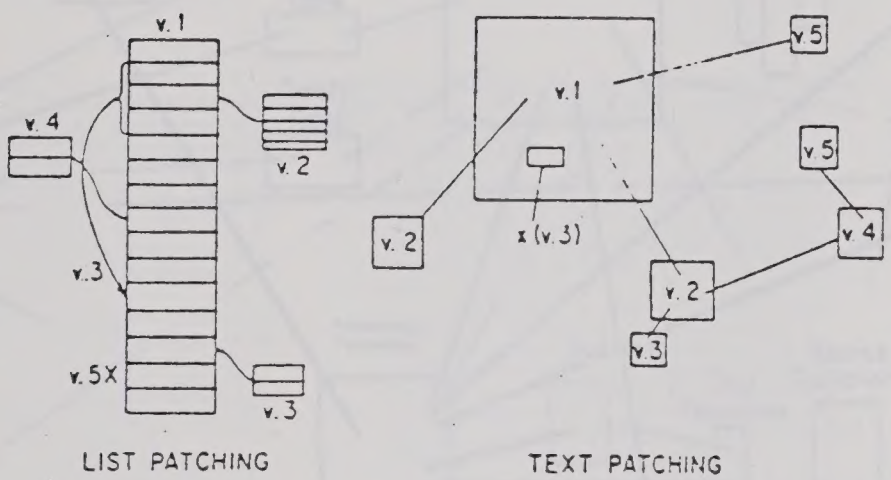

FIGURE 2—Spinoff of variants: extra versions need little space.

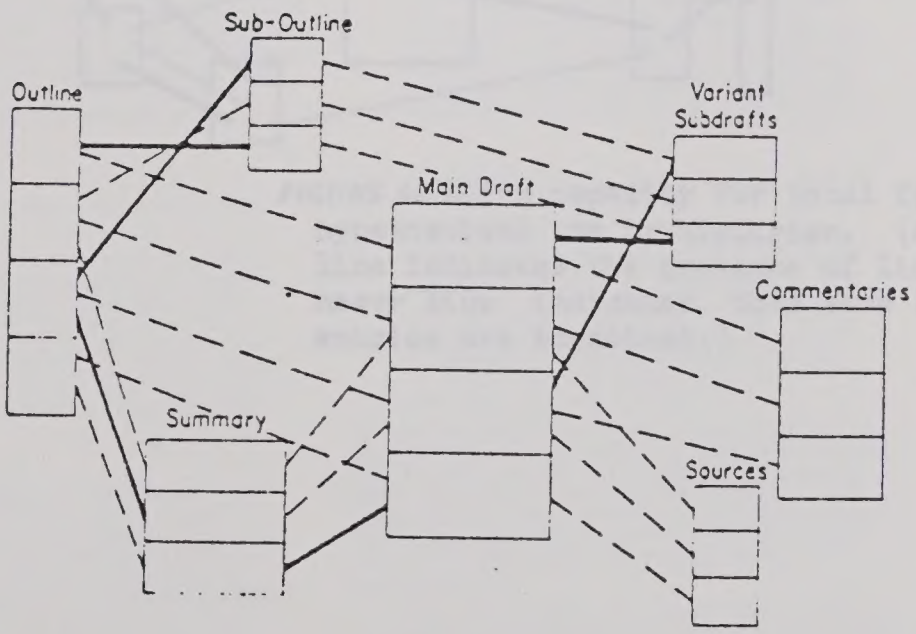

FIGURE 3— All levels of documentation may be contained in the ELF. (Heavy lines indicate that linked entries are identical.)

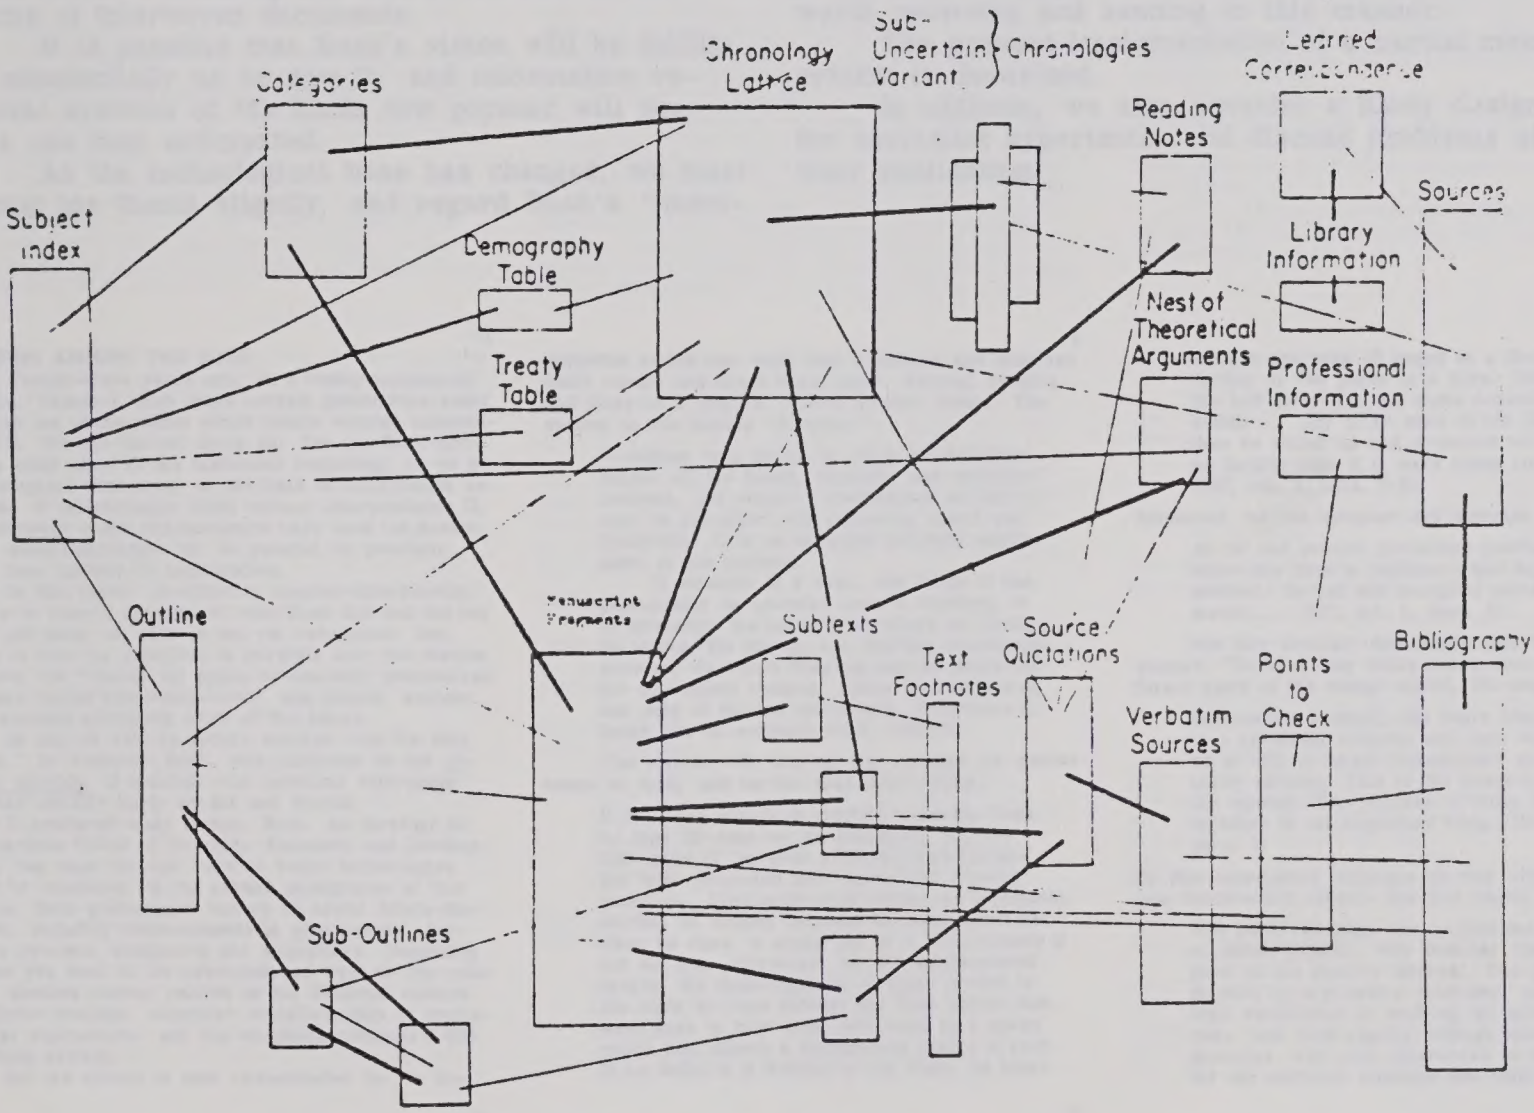

FIGURE 4– ELF's capacity for total filing: hypothetical use by historian. (A thin line indicates the presence of links; a heavy line indicates that some linked entries are identical.)

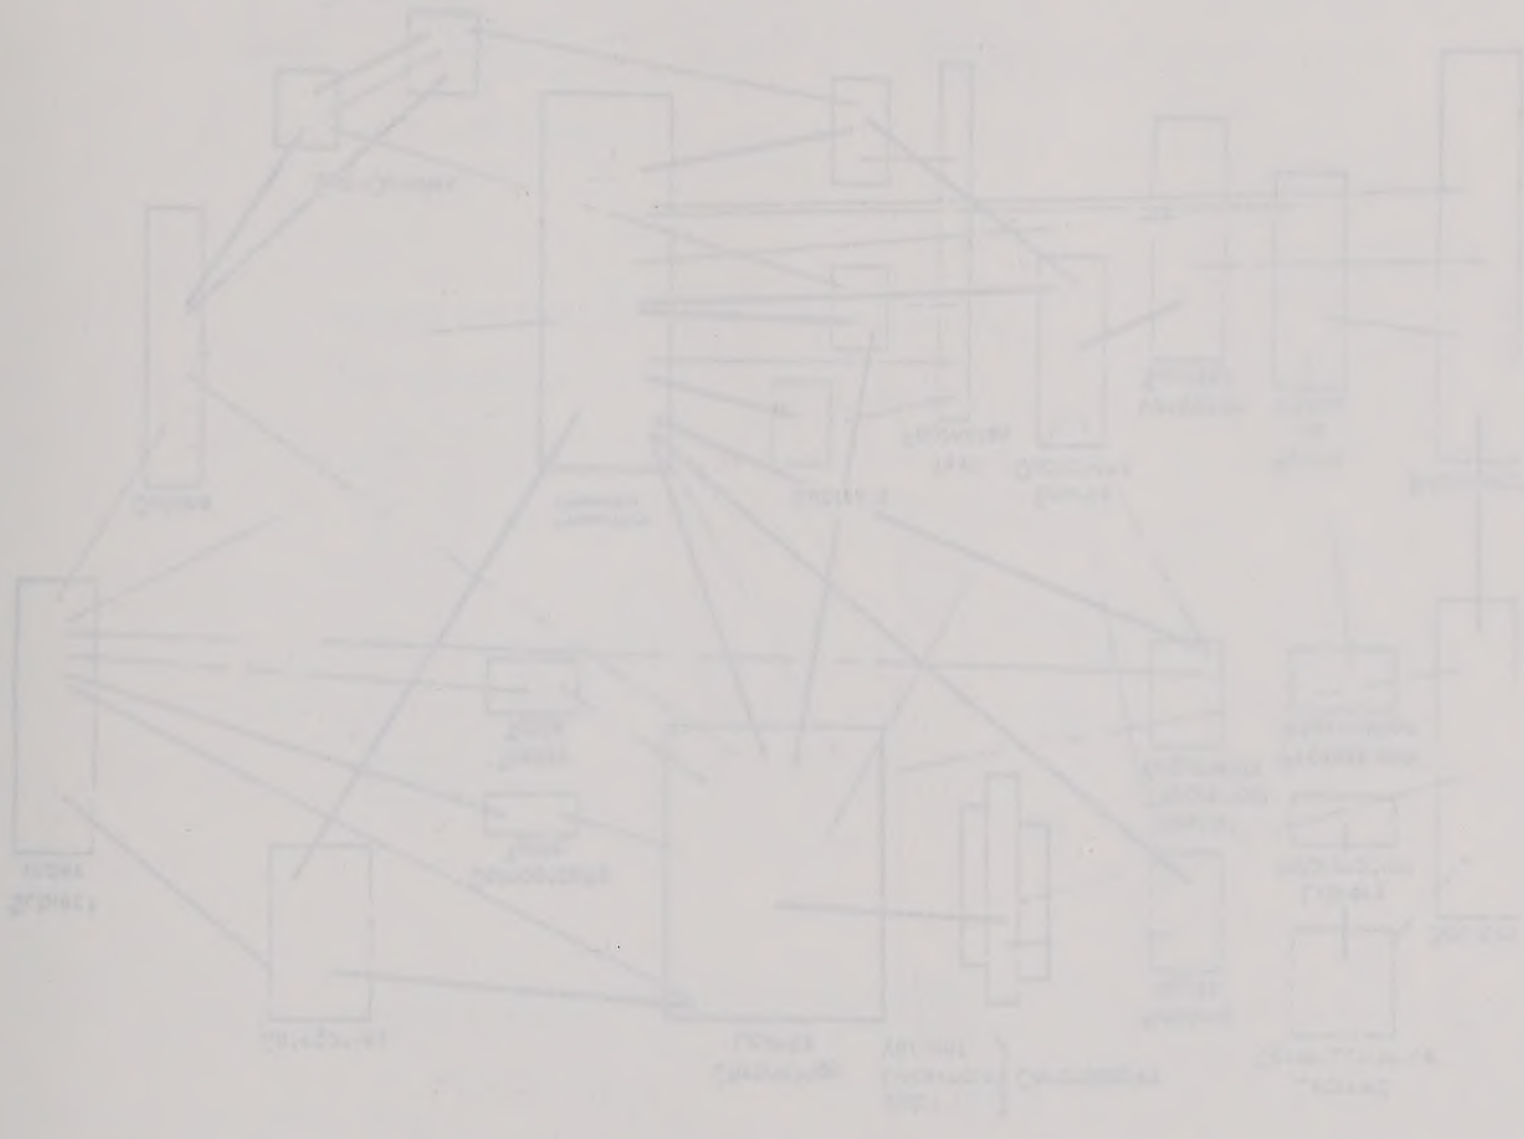

(50)

ABSTRACT.

AS WE WILL THINK. Theodor H. Nelson. 458 W. 20th Street, New York, N. Y. 10011.

Bush was right. His famous article is, however, generally misinterpreted, for it has little to do with "information retrieval" as prosecuted today. Bush rejected indexing and discussed instead new forms of interwoven documents.

It is possible that Bush's vision will be fulfilled substantially as he saw it, and information retrieval systems of the kinds now popular will see less use than anticipated.

As the technological base has changed, we must recast his thesis slightly, and regard Bush's "memex" as three things: the personal editing and file console; a digital feeder network for the delivery of documents in full-text digital form; and new types of documents, or hypertexts, which are especially worth receiving and sending in this manner.

The present implementation of a partial memex system is described.

In addition, we also consider a likely design for specialist hypertexts, and discuss problems of their publication.

1

BEATING AROUND THE BUSH

Twenty-three years ago, in a widely acclaimed article, Vannevar Bush made certain predictions about the way we of the future would handle written information (1). We are not yet doing so. Yet the Bush article is often cited as the historical beginning, or as a technological watershed, of the field of information retrieval. It is frequently cited without interpretation (2, 3). Although some commentators have said its predictions were improbable (4), in general its precepts have been ignored by acclamation.

In this paper, an effort in counter-discipleship, I hope to remind readers of what Bush did and did not say, and point out what is not yet recognized: that much of what he predicted is possible now; the memex is here; the "trails" he spoke of--suitably generalized, and now called hypertexts--may, and should, become the principal publishing form of the future.

In July of 1945 an article entitled "As We May Think," by Vannevar Bush, was published in the Atlantic Monthly. It bristled with technical references but was actually fairly candid and simple.

It predicted many things. Bush, as director of the wartime Office of Scientific Research and Development, had seen the new ways in which technologies could be combined. In the urbane paragraphs of this article, Bush predicted a variety of useful future machines, including improvements in photography, facsimile systems, computers and miscellany. Depending on how you read it, he predicted, as well as you could hope, devices closely related to the Polaroid camera, the Xerox machine, computer transformation of mathematical expressions, and the telephone company's ESS switching system.

But the article is best remembered for its description of the new ways that scientists and scholars could handle and share their ideas, writing, reading and filing in a magical system at their desks. The system is the famous "memex."

2

> A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.
>
> It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works. On the top are slanting translucent screens, on which material can be projected for convenient reading. There is a keyboard, and sets of buttons and levers. Otherwise it looks like an ordinary desk. (106-7)

The memex will hold all the writings its master wants to read, and he can read them easily.

> If the user wishes to consult a certain book, he taps its code on the keyboard, and the title page of the book promptly appears befor him, projected onto one of his viewing positions. Frequently-used codes are mnemonic, so that he seldom consults his code book; but when he does, a single tap of a key projects it for his use. Moreover, he has supplemental levers. On deflecting one of these levers to the right he runs through the book before him, each page in turn being projected at a speed which just allows a recognizing glance at each. If he deflects it further to the right, he steps through the book 10 pages at a time; still further at 100 pages at a time. Deflection to the left gives him the same control backwards.... Any given book of his library can thus be called up and consulted with far greater facility than if it were taken from a shelf. (107, col. 1, para. 5-6)

3

Moreover, he can compare and annotate them.

> As he has several projection positions, he can leave one item in position while he calls up another. He can add marginal notes and comments,... (107, col. 1, para. 8)

Not only ordinary documents need be held in the memex. The user may make connections between different parts of the things stored. He does this by

> associative indexing, the basic idea of which is a provision whereby any item may be caused at will to select immediately and automatically another. This is the essential feature of the memex. The process of tying two items together is the important thing. (107, col. 2, para. 1)

By this associative technique he may create "trails," new documentary objects that are useful in new ways.

> The patent attorney has on call the millions of issued patents, with familiar trails to every point of his client's interest. The physician, puzzled by a patient's reactions, strikes the trail established in studying an earlier similar case, and runs rapidly through analogous case histories, with side references to the classics for the pertinent anatomy and histology. (108, col. 1, para. 1)

4

These new structures, or trails, may be taken and given to other people.

> Tapping a few keys projects the head of the trail. A lever runs through it at will, stopping at interesting items, going off on side excursions. It is an interesting trail, pertinent to the discussion. So he sets a reproducer in action,... and passes it to his friend for insertion in his own memex, there to be linked into the more general trail. (107, col. 2, para. 5)

And they may be published.

> Wholly new forms of encyclopedia will appear, ready-made with a mesh of associative trails running through them, ready to be dropped into the memex and there amplified.... There is a new profession of trail blazers, those who find delight in the task of establishing useful trails through the enormous mass of the common record. The inheritance from the master becomes, not only his additions to the world's record, but for his disciples the entire scaffolding by which they were erected. (108, col. 1, para. 2)

It is strange that "As We May Think" has been taken so to heart in the field of information retrieval, since it runs counter to virtually all work being pursued under the name of information retrieval today. Such systems are principally concerned either with indexing conventional documents by content, or with somehow representing that content in a way that can be mechanically searched and deciphered.

5

This is paradoxical. On the one hand, Bush did not think well of indexing.

> The real heart of the matter of selection, however, goes deeper than a lag in the adoption of mechanisms by libraries, or a lack of development of devices for their use. Our ineptitude in getting at the record is largely caused by the artificiality of systems of indexing. ... [between documents] one has to emerge from the system and re-enter on a new path. (106, col. 2, para. 3)

On the other hand, Bush merely hinted about the use of structured-data representations and calculi in storing ideas (105, col. 2, para. 3), and did not plainly relate them to his main exposition. While we might argue scripture about such matters, the fact is that Bush's most extensive concern has had few successors in the field called "information retrieval."

TRANSPOSITION

The memex was to be a single screen console for handling the user's notes, writing and correspondence, for reading books and other writings created by others, and for creating new associative text structures, which may in their turn be read and distributed. All this I take to be the heart of Bush's prediction. This will happen. Such systems exist; they are approaching cost feasibility; and the world is readier than it thinks.

Bush's machine will not, of course, be built exactly as he foresaw. The complete description, which I omitted, involves microfilm cassettes, a photographic copying plate for adding new images to the microfilm file, and a telautograph stylus. Other

6

machines he describes, such as the forehead camera and the direct-dictation typing machine, might or might not have been coupled to it as well. In the revised version of the article (5) his emphasis shifted to videotape. These impedimenta we ignore. There exist microfilm cassettes, copiers and so on, but the hardware ready to support a memex-class system will be something else.

The system will be built from existing computer equipment and peripherals. Physically it will be a computer display, with a keyboard, at the user's desk; a support computer system (at the desk or elsewhere) for handling various technical chores; and a library network of digital feeder machines. The written materials, when not shown on the screen, will be stored and sent digitally, as telegraphic symbols. They will be sent back and forth among these systems automatically, as programmed in the various devices. The trails, or associative text structures--more generally called "hypertexts"--will be stored in coded form, along with the more conventional documents in the various devices. The user will be billed automatically for the services and the delivery of copyrighted materials. The publisher, who maintains these copyrighted materials in the feeder machines, will be duly paid for their use.

Prototype units exist now. Appropriate console hardware can be purchased now for about $15,000.

Supply systems, however, are not quite ready. The best supply system may be a special-purpose computer or a general-purpose time-sharing system; which is better is not clear. While costs of either are presently in the hundreds of thousands, they are coming down fast, and the use of a well-tailored system by many people at once should bring the cost of such ser-

(51)

vice down considerably. To name a figure arbitrarily, let us say that a service cost of $100 a month per user (exclusive of telephone lines and copyright) would be sufficiently low to draw many users. Such service at such a cost will surely be generally available between one and five years from now.

Various preoccupations have delayed us psychologically. We do not need direct dictation, optical scanning or the availability of vast libraries for such systems to be immediately practical and important to us.

THE CONSOLE

We are speaking of a single console to handle notes, writing, much correspondence, much reading, and the creation of new kinds of texts. On it the user must be able to view, edit, file, and otherwise manipulate.

Let me now describe an existing system. This system was created for the purpose of experimenting with various kinds of text handling, and to aid in the production editing of documents which require team study and revision.

The Brown University hypertext editing system, created by Andries Van Dam, Walter Gross and me, has a number of facilities which turn out to be those of Bush's memex. Conscious as we were of the Bush article, we were not paying attention to these parallels during the design deliberations.

Many of these features exist in other text handling systems, notably that of Douglas C. Engelbart at Stanford Research Institute (6). The purpose of this description is to show parallelisms between memex and this general type of system, not to distinguish this system from its relatives.

I will speak of the Brown hypertext system simply as "the system," "the current system" or the like, to distinguish it from the memex. As its implementation is not yet complete, this description applies to the present specifications and not only the parts that are working.

The user sits at a foot-square screen with a desk-top, typewriter keyboard, and a special set of pushbuttons. (This is the IBM 2250, but lesser facilities, in this setting, would be capable of comparable service.)

A. Text handling.

First let me briefly describe the text-handling and text-editing facilities of the system. While this is a sidetrack, since Bush did not discuss editorial work, it is a basic aspect of the system. (This is a great advantage of digital over pictorial storage: the system may manipulate the words letter by letter, rather than as a single image.)

The basic text-handling features of the system are the presentation of text materials on a screen; the ability to command basic editing operations by simple manipulations; the ability to make these editorial changes tentatively, on copies or alternate versions of your material; and the option of having your actions recorded automatically in a cumulative editorial log, in case it is later necessary to retract any of them.

The system is generally geared to tentative and thoughtful operation. The alternative versions and editorial log are the strongest examples, but there are others. An editing command, once enacted on the screen, must still be accepted by the user before it actually occurs in the file. Another example: a section of text being tried in a new place is shown with a

number of spaces to the left of its beginning and the right of its end, so you know its limits within the new setting.

B. Filing and manipulation.

Various techniques must exist for manipulating the file in such a system.

The user of Bush's memex called its contents by means of a "code book" (107 passim)--but this too was in the system. From the code book the user was to choose contents for viewing by looking up and keying in their listed codes--some mnemonic, some shortened over use.

In the current system it is possible for the user to draw something to the screen by typing its user-assigned name, or picking the name from a screen menu with the light pen. He may do this not just with whole units, but also summoning instances of things to be repetitively copied, and different versions--separate copies--of documents to be tentatively edited.

In the memex, the user could skip through a document at adjustable rates (107, col. 1, para. 5). In the current system it is possible to step forward through the text by an adjustable number of lines at a time.

The memex user looked at several documents simultaneously through several screens or panels of the same screen (107, col. 1, para. 6). This will be possible for up to three documents on the current system.

The memex user created annotations by hand, or links between the things that were being viewed simultaneously.

The user of the current system has two possible types of annotation he can create. The first, presently called the "tag," is a marker in the text to which a short note or explanation is attached. This may be typed in at the keyboard. From the tag there is no place to go except back. The second connector, called in this system the "link," joins two items of text. Either item may be anywhere in the user's file. A link may also be made to a nonexistent text item, which the user may then put in by the keyboard.

The current system assumes no structure, but creates links between text sections regardless of whether or not they are parts of the same unit or otherwise related.

In addition, the memex described in the revised version of Bush's article kept a log of manipulations by the user ("Memex Revisited," 95-6). The current system will also possess this capability.

Design problems arise in the richer operations, such as creating connective structures. It is taken for granted that the console must be easy to use. That is no design problem for a small set of operations, such as text editing. But the design of the overall file handling elements and actions is more complicated. The problem is providing conceptual unification for the system's filing structure and conventions.

For example, in the current system, the number of basic file connectors is held to two.

The basic idea is that the tag (signified by an @ in the text) is a machine-scannable annotation, the link (signified by an * in the text) a basic jump from one item to another. Tags may be of various types, and each type may be collected, listed and searched by the machine. Links, marked by an asterisk, may

also have a textual "explainer" telling what kind of a jump is available, or anything else the hypertext's creator wants to attach to the link.

This dichotomy, while simple, should be sufficient to support real hypertext experimentation.

Complex couplings are presently not defined for the current system. However, the textual labelling of links is adequate for various possible forms of discrete-jump hypertexts, including Bush trails, and may in principle be extended to computer responsibility for link behaviors and complex coupling maintenance.

HYPERTEXTS

While Bush's term, "trail," represents a very useful concept, we must generalize it. Bush's interest in microfilm led to his idea of the trail having a sequence.

By "trail," Bush appears to have meant a sequence of documents, document excerpts, and comments upon them.

> For [the user] runs through an encyclopedia, finds an interesting but sketchy article, leaves it projected. Next, in a history, he finds another pertinent item, and ties the two together. Thus he goes, building a trail of many items. (107, col. 2, para. 4).

This sequence would be established by making paired couplings.

> When the user is building a trail, he names it, inserts the name in his code book, and taps it out on his keyboard. Before him are the two items to be joined projected onto adjacent viewing positions. ... The user taps a single key, and the items are permanently

joined. (107, col. 2, para. 2)

Bush mentions two other types of trail. One is the "side trail," branching out from a main trail sequence.

> Occasionally he inserts a comment of his own, either linking it into the main trail or joining it by a side trail to a particular item. (107, col. 2, para. 4)

The other type of trail is the "skip trail," a subset of a main trail sequence that contains the highlights.

> The historian, with a vast chronological account of a people, parallels it with a skip trail which stops only on the salient items, and can follow at any time contemporary trails which lead him all over civilization at a particular epoch. (108, col. 1, para. 1)

In Bush's trails, the user had no choices to make as he moved through the sequence of items, except at an intersection of trails. With computer storage, however, no sequence need by imposed on the material; and, instead of simply storing materials in their order of arrival or of being noticed, it will be possible to create overall structures of greater useful complexity. These may have, for instance, patterns of branches in various directions. Such non-sequential or complex text structures we may call "hypertexts." (7)

"Hypertext" is the generic term; there are reasons, for which there is no room here, to rule out such other candidate terms as "branching text," "graph-structured text," "complex text" and "tree text."

The best current definition of hypertext, over

quite a broad range of types, is "text structure that cannot be conveniently printed." This is not very specific or profound, but it fits best.

As Bush pointed out in his own terms, we think in hypertext (106, col. 2). We have been speaking hypertext all our lives and never known it. It is usually only in writing that we must pick thoughts up and irrelevantly put them down in the sequence demanded by the printed word. Writing is a process of making the tree of thought into a picket fence.

Hypertext structures are varied. For instance, they may be free-branching with only one type of link and backing up; they may have modal links with different meanings in a free structure; or have modal links and repetitive structure.

Discrete-jump hypertexts are not the only kinds. There is, for example, "stretchtext." This is continuously variable text which never leaves the screen, but changes by small increments on user demand, growing longer and more detailed by a few words at a time, as required. Other continuous types are possible.

As items may be coupled, whole hypertexts may be coupled into books, or one another. An example of the first would be annotations coupling into the Bible. Such multi-couplings involve bundles of pointers between the texts, possibly with type codes or annotations. They may also involve alternative versions, which there is no room to discuss here. The structuring of these coupler types is a continuing design task for hypertext systems.

The creator of hypertexts may allow the user various options of jumping or branching. These options can lead the user to further reading in any pattern the author wants to make available to him. The only constraints on the author are usefulness, clarity, and artfulness.

There must, of course, be ways the reader may see, and choose among, possible branches from an item in the text. This problem was implicit in Bush's treatment. Since "any item can be joined into numerous trails" (107 col. 2 para. 3), there would have had to be some way of showing the user these options and their meaning. This is the case in general, and a standing aspect of hypertext system design.

Hypertexts may be casual rough notes, as described in Bush's extended example of the Turkish bow and arrow; or they may at the other extreme be finished units, editorially completed and organized. Such finished units would have many of the same properties as ordinary writing: intentional assembly, attempted clarity, and expository structure of enumerable "points," and an overall comprehensible pattern whose interrelated parts may be in some way remembered or visualized. Finally, the concept of authorship applies to hypertexts as much as it does to an ordinary book or article.

As with ordinary texts, too, the editorial properties and "feel" of hypertexts may be quite distinct and varied. For hypertexts these are of course largely unexplored. It is also very hard to anticipate their possible administrative and social settings, and this will greatly affect their character and the modes of their use.

THE TRANSMISSION NETWORK

A general transmission network will carry requested documents from libraries to users, new documents from users to libraries, and communications and documents between users.

The network will consist of several computers or

computer-like objects. In the user's own unit is digital logic, and possibly a small computer; this unit is serviced or managed by a computer which stores the user's files and communicates with the library network. The user's requests for documents that are not available locally go out to a library network. These requests are sorted out to the appropriate repository machine; the repository machine returns the document and a bill or fee schedule.

Various fees are logged up to the user. These will include various basic costs, such as membership in the system, rental of the terminal and hookup. Additional fees may include logged-in time, per-usage costs of various facilities (such as average memory area occupied and quantities of text moved), storage charges for materials kept locally, and royalties to copyright holders. It should be noted that various grades of service may also exist, in which the user gets faster service by claiming larger buffers and higher priorities, and pays for the privilege.

Although this may sound like a formidable prospect, in general and with polishing there is reason to hope that the real costs of such a system will compete favorably with the real costs of the forms of publication and libraries we now employ. (Of course, in such "real cost" we must include the library services supplied "free" by various levels of government, including municipal libraries, grants to universities and the indirect subsidy of publishers by low postal rates. It is not unthinkable that similar encouragement will come about for this form of publishing and libraryship.)

Various technical design issues exist. These involve the feeder computers and their forms of memory. These hierarchies of memory are fairly clear.

They will generally include disk (for working areas and directories), magnetic card or data cell (for the corpus), and magnetic tape (for rarely-needed materials and safety copies).

A more difficult question is, what should the feeder computers be? Their job in this system is the lookup and shovelling of text, plus bookkeeping. One school of thought holds that a true general-purpose time-sharing system is necessary; another, that the correct machine is a dedicated computer with rich interrupts and comparatively little arithmetic capacity. The third school would point out the special character of the work and lean toward special designs and special tradeoffs, which could be anything from associative memory to the use of delay-line machines.

Similar issues exist for memory software and directory systems. There are complex technical problems of index and search techniques, and methods of their cross-tabulation. But they can be handled in some way or other.

A warning is necessary here, however. This area of console support is the area where things are not yet ready. The prediction of economic feasibility in five years, an eon in the computer world, is not the same as feasibility now. By devoting a whole computer, disk and tape to each user, the problem of console support can now be solved, in a manner of speaking; but the general problem of interleaved I/O and file management, with the efficient sharing of facilities, is another problem entirely, and the one that must be solved to make this whole thing go.

PUBLICATION: REDESIGN OF THE TECHNICAL LITERATURE

Bush regarded his new text structures as trans-

missible between individuals, and publishable. The same is true of hypertext units, the generalized form of trails. I think it likely that once such systems are available, the creation of branching and complex text will become recognized as far more natural than the structures in which we now must write.

This will all follow naturally from the existence of consoles which permit multiple couplings between texts. Having created for personal use a hyper-document on one's console, it will seem only reasonable to press a button passing this on to a colleague in its hyper-form, without chopping and aligning it into conventional writing.

Various interesting possibilities follow. Private "journals" in a field may be started among co-workers merely by the pooling of their hyper-documents. The rental of memory space on magnetic cards is inconsequential next to what have been the costs of printing and mailing.

When professional and technical societies become interested in sponsoring hyper-publication, one of the most straightforward ways to begin would be with the creation of society-sponsored review articles. These could be like the ordinary review articles sponsored by such societies, save that the review article would open directly into the various materials it was reviewing, and footnotes could be more extensive and slanted to different categories of reader. (Figure 1.)

The next step is, of course, the creation of hyper-magazines or journals under the sponsorship and supervision of professional societies. Here the problem of organization would seem to become thorny. But this collection could be much like the journals of today, except for the direct availability of previous liter-

ature, working papers and various odds and ends. (Figure 2.) It should be observed that any of the "documents" noted in the illustration can themselves be hypertexts.

In this magazine, an arbitrary conjecture, all material of the past year is considered "recent" and embraced in a common lookup structure. New material enters the collection surreptitiously, at whatever time of day or night the editors release it; material one year old is formally expelled to a different file on the first of each month. (It may be just as accessible, but its nominal status changes, rather like that of a book taken off the "recent acquisitions" list of a library.) This magazine would hold most of what people were talking about in a field, and it would all be right there.

There are numerous technical complications to hypertext publishing. There is no room here to discuss the more esoteric technical ones, such as facilities, billing and copyright conventions, or possible techniques of encryption and validation to prevent pirating of works and to authenticate expeditionary versions.

HOW TO BEGIN

How shall it begin?

Those contemplating massive retrieval systems commonly presume that they must begin with some massive corpus all accessible. The Library of Congress is often mentioned. Even Bush supposes regretfully (in the revised article, p. 100) that the personal system waits on large public establishments being automated first.

I do not believe this is so. It will be practicable and of considerable interest to begin on a small scale,

having no grand corpus available. The grand corpus will come soon enough, as requests emerge. (We have a precedent: the prowess of University Microfilms, Inc. in rendering texts available to scholars in microfilm.)

The way to begin is to furnish supported consoles to small communities of users: key members of a "small" discipline, or specialists among whose work there is close connection. Suitable groups might be "early Egyptologists," or just plain "everybody at Woods Hole."

Such communicants, having been assured as well as possible of privacy and fail-safe design, will be encouraged to use the consoles fully. From the outset they may keep all their notes, manuscripts, articles, and copies of outgoing correspondence, on the system.

The rest will follow. I am fairly sure of the predictions so far, at least in broad outline, but I am just as sure that the first generation of hypertext users will invent twice as much as has been descried and described so far.

Who will support these beginnings? We have a choice, at the outset, of universities, publishing companies, computer companies (including service bureaux), research organizations or the government. Any of these might take such an initiative. Though such an initiative would seem severally unthinkable, it somehow seems collectively plausible and, of course, historically inevitable. If you believe in manifest destiny.

References

1. Vannevar Bush, "As We May Think." The Atlantic Monthly 176:1 (July, 1945), 101-108.

References to "Bush" cite this article unless

further specified, and all unexplained numbers and paginations in the text refer to it. Because the article is so tightly written and our interest in it so close, paginations are given to the column and paragraph.

2. Allen Kent, Textbook on Mechanized Information Retrieval, second edition. Wiley, 1966; pp. 7-9.
3. Joseph Becker and Robert M. Hayes, Information Storage and Retrieval: Tools, Elements, Theories. Wiley, 1963; p. 40.
4. John H. Wilson, "As We May Have Thought." In Progress in Information Science and Technology: Proceedings of the American Documentation Institute 1966 Annual Meeting. Adrianne Press, 1966, pp. 117-122.
5. Vannevar Bush, "Memex Revisited." In Vannevar Bush, Science Is Not Enough. William Morrow and Company, Inc., 1967, 75-101.
6. Douglas C. Engelbart, "Augmenting Human Intellect: A Conceptual Framework." AFOSR-3223 Summary Report, Stanford Research Institute, Menlo Park, California, October 1962.

This expresses the philosophy, rather than the results, of the continuing project.

7. Theodor H. Nelson, "Getting It Out of Our System." In George Schechter (ed.), Information Retrieval: A Critical Review. Thompson Book Company, 1967; pp. 191-210.

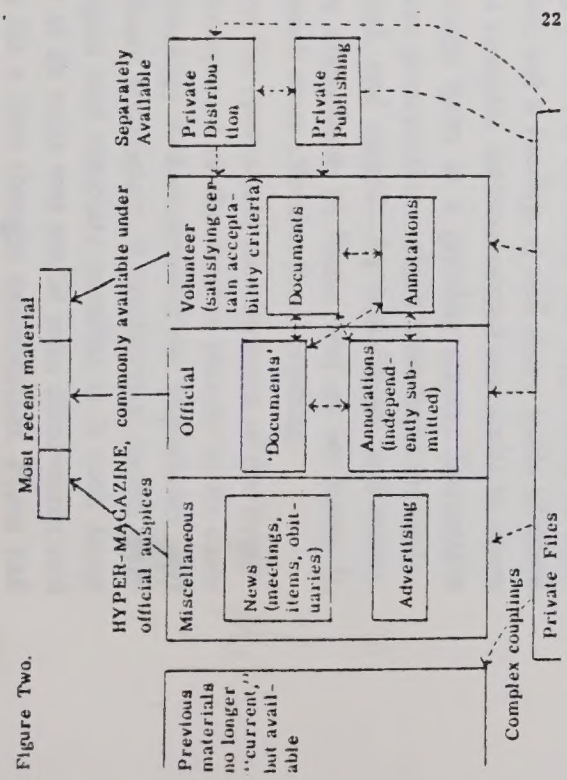

Figure Two.

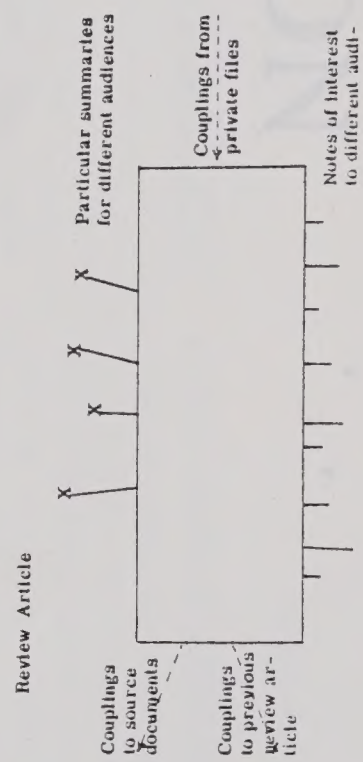
Particular summaries for different audiences
Couplings from private files
Notes of interest to different audiences
Review Article
Couplings to source documents
Couplings to previous review article

Figure One.

INFORMATION RETRIEVAL

A Critical View

Edited by GEORGE SCHECTER

Based on Third Annual Colloquium
on Information Retrieval
May 12-13, 1966, Philadelphia, Pennsylvania

Sponsored by

Special Interest Group on Information Retrieval, Association for Computing Machinery

Delaware Valley Chapter, Association for Computing Machinery

Institute of Electrical and Electronic Engineers Computer Group, Philadelphia Section

United States Army Frankford Arsenal

Moore School of Electrical Engineering, University of Pennsylvania

1967

THOMPSON BOOK COMPANY
Washington, D. C.

ACADEMIC PRESS
London

REPRINTED BY PERMISSION

GETTING IT OUT OF OUR SYSTEM

THEODOR H. NELSON
Harcourt, Brace & World, Inc.
New York, New York

SUMMARY

Specialist access and muster of information—usually the underlying task of "information retrieval"—may not be best accomplished either by indexing techniques (document retrieval) or queriable information networks (content retrieval). Digital text storage and display make possible the creation of at least one potentially powerful *medium*: the hypertext, or nonlinear text system. As a medium, not a facility, it will share traits of book and film. Machines will present complex interconnected text units to users who will weave their way through them at a glowing screen. Successive user choices will evoke successive presentations which have been created—written and structurally designed, and signed—by individuals and teams. Departing completely from the convention of linear text as we know it, the hypertext will probably be much more effective for everything.

Several parallels exist between the (potential) development of the hypertext medium and the actual development of the motion picture in its early years: the inventors of the technology underestimated the importance and generality of the resulting medium; techniques of predecessor media were imported to it, and the techniques and structure peculiar to it took some years to develop; the character of its units was not at once discovered; and credit to a work's creators was for a time thought unimportant. Films had

RETRIEVAL
INFORMATION

one special advantage: a standardized technical context existed to make the exchange of units simple.

INTRODUCTION

What do we want to get out of an information system? There is a great need to turn the new machines—computers, mass memories, and dynamic displays—into information systems, or, to be more precise, information storage and supply systems (as distinct from systems for information processing.) But how are they to work, and what are they to do?

What do we want to get out of an information system? The way we answer this rhetorical question probably settles the major aspects of the system we design. The hard-nosed answer is "Documents." The dreamy-eyed answer is "Pure information." The hesitant answer is to ask another question. "What sorts of things is the system for?" Let us probe this subquestion neutrally. If we try to think in terms that avoid the stereotypes and commitments of information retrieval, perhaps we can find a new way of seeing the problem context, and better understand the conventional solutions for what they are and do.

What do we want to get out of an information system? The information problem I will confront here is the overall question of how to keep specialists informed and updated. To give it an acronym, I would venture that existing systems are *S*pecialist *A*ccess and *M*uster of *I*nformation (SAMI for short). To speak of "access and muster" avoids the connotations of "retrieval"—that the information has been lost, prodigal, or shot down in the bushes. (We will exclude other possible meanings of the term "information retrieval," considering library inventory systems to be a temporary expedient, and management information to be a job for real-time system programming.) I will propose that SAMI is best served, not by document retrieval or content retrieval, each in its way more conventional than it looks, but by something else again. This paper is an elaborate conjecture which I hope to give verisimilitude by an extended historical analogy.

"INFORMATION RETRIEVAL"

Information retrieval in the present sense now means two things. One, document retrieval, uses standardized indexing techniques to draw desired classes of unit documents from a central pool. The other, content retrieval, holds a body of integrated knowledge to which the user may address questions. (For a unified treatment of the two types, see Licklider.[1]) Each is important.

A. DOCUMENT RETRIEVAL

Let us examine the first meaning of information retrieval, document retrieval (DR), the finding of documents having certain indexed descriptions. In document retrieval we index a collection of documents by particular terms or aspects, and later only get things back *through* those terms. This approach retains the unit of the discrete "document," taking it as given. The documents themselves are not expected to have any provision for the system, or to have been created with any relation to the system in mind, but are produced by the world according to external causes. The document is integrally sacrosanct, its content and form untouched by the system. No tampering, excerpting or post-editing of the documents is permissible, no matter how useful it might be.

The problem of the system then becomes finding "the right" documents, which the user recognizes by criteria of his own. The match between the user's private criteria and the available indexing is indicated by the number of "hits"—documents the user is pleased with.

Studies of library use have indicated that not very much library use could be aided by conventional document retrieval systems as currently construed. This would be all right if we could be sure they were good for something else.

That is one criticism of document retrieval: it is not yet clear that document-retrieval systems will do what the users need to have done, even in some ordinary sense, and before we start reforming the users.

The second criticism of document retrieval is based on the nature of the indexing, or categorization, process. (Indexing is categorization, even when it becomes complicated with the relational tagging of roles and links.)

The best categories, and indices, decay. In principle any indexing categories will progressively decline in relevance. It is quite reasonable to expect that as emphasis and interests change in a field, many or most indexing categories for information retrieval will gradually fall apart. A cursory examination of the Dewey Decimal System should confirm this point. If the structure of a particular field strays from the way the original indexing saw it, we will look in the wrong place for the information that was put away earlier.

B. CONTENT RETRIEVAL

The second sense of information retrieval now being used is content retrieval (CR). Content retrieval is often called "fact retrieval"; a system for content retrieval is sometimes called the QAS, or question-answering system. (An excellent source on CR is Simmons.[2])

Such systems can be seen as having two parts; some overall method of storing a complex network of interrelated information, such as that of Marill,[3] and a program mechanism for following through its different parts to find something out. Network and query programs are intertwined; the design of one has consequences for the other. However, for the present purposes we may call such a system the "queriable information network," since the information network is the basic part; in principle a way can presumably be found to interrogate any adequate logistic system.

Content retrieval discards the unit of the document. The total contents of a queriable information network need not be divided into units at all, except those described in the information and those required to represent it.

Such systems offer great promise in a number of ways. However, they do have present and future drawbacks. Whether such systems can replace significant amounts of library use cannot be ascertained. Second, preparing information to go into the queriable information network may be an extremely difficult, touchy and expensive business, at least as things look now. For some systems, coders may require training in mathematical notations of great complexity. Third, conversing with such a machine will on some systems usually, and in the most lucid systems sometimes,[4] require formulating questions in arcane languages or dialects. Either these can be known only to specialists, or their use by the laity may lead to mistakes.

The major problem, though, and the real long-term drawback, is analogous to the decay of categories and indices in DR. Content retrieval assumes that all the information in the network is *consistent and true*, and has no inconvenient gaps. Like any overview, CR must ignore uncertainties and lacunae, and (as "knowledge" is revised and interests shift) it may leave only the tintype of a certain moment's concerns. As scientific and other beliefs are constantly changing, "knowledge" must frequently be readjusted or rearranged.

Actually, the idea of "knowledge" is misleading in this context. What we regard as knowledge includes various contradictory sets of interwoven propositions, clusters of argument, and discrepancies of faith, taste and emphasis in the literature. There is usually disagreement in every field of knowledge. This means that, in a queriable information network, different and mutually contradictory theories really ought to be able to coexist within the system.

Moreover, we may surmise that in most fields of knowledge some noncontroversial information is incorrect, but we don't know which. It follows that unless the information network can be modified as systems of information are modified, the information network will be left behind. While evolutionary and patching modifications might be arranged to cope with the representation of changing views, this expands the problem significantly. (For more on such evolutionary issues and difficulties, see Nelson.[5])

MANIFESTO

I would like to make the claim that document and content retrieval systems are not the principal nor the basic use of computers for handling ideas. I am not claiming that such systems are useless or improper. Rather, I think they will be helpful accessories to the proper, natural and true use of computers for handling the written word.

Systems thinkers have proposed a great variety of possible arrangements for ordering man's thoughts and aiding his mind with the computer. I would like to propose that this diversity of possible facilities and functions is in fact comprisable into an *organized text medium*, to be created and used by people and for people. This organized medium will be that of the text structure created expressly for on-line display by the computer. Such a text will be read from an illuminated screen, the cathode-ray display; it will respond or branch upon actions by the user. It will be a succession of displays, largely textual, that come and go according to his actions. Such a congeries of interconnected text may incorporate and have available many chunks of literature on a given subject, or many subjects, purposely assembled and woven together by authors and editors. The name for this new medium is hypertext.[6]

Hypertext is the combination of natural-language text with the computer's capacities for interactive, branching or dynamic display, *when explicitly used as a medium*. Or, to define it more broadly, "hypertext" is the generic term for any text which cannot be printed (or printed conveniently) on a conventional page, or used conveniently when bound between conventional covers. "Non-linear text" might be a fair approximation. Hypertext may differ from ordinary text in its sequencing (it may branch into trees and networks), its organization (it may have multiple levels of summary and detail), its mode of presentation (it may contain moving or manipulable illustrations, moving or flashing typography), and so on. At last, with the appearance of the computer display and appropriate back-up systems, such texts are possible and practical.

We will store natural-language texts in digital devices for dynamic display, and arrange them into new kinds of units. This means discarding customary document boundaries, and organizing the materials in a complex and non-linear fashion. The user will read from a screen, and as he goes he will make choices that call connected materials into view. Hypertext will be a medium, not a facility, and the institutions of authorship, citation,

academic commentary and argument will be preserved and incorporated in an orderly fashion.

Thousands (or millions) of paragraphs of source material may be available directly to the inquiring reader who wishes further detail, or a certain person's comment, or a specific idea or sentence. Moreover, the *original context* of such items may be summoned, and such contexts—articles, books, or other hypertext units—themselves examined. Thus the reader in search of understanding may continually uncover new elements of interest, often of growing relevance to his hunches, his confusions and his unphrased questions. He can find explanations, ideas, or whole new areas that are to his interest. He may study them through further excerpts, reviews, summaries, or even conventional reading.

The reader will explore, light pen in hand, browsing or studying closely, choosing paths he prefers, criss-crossing these paths on summary levels, stepping from section to section and stopping for deeper work in units of interest. He will come to understand the whole, or a general section, like a walker exploring a city.

Hypertexts will require no special training to use, and rather little special training to create, though like works in other media they can be made either well or poorly. To use, or read, a hypertext, you will need mainly to sit at a screen, read, and indicate choices by pushbutton or light pen. To create a hypertext, you will type in new text and indicate the new relational connections (and their operational meaning) to the system. Where you wish to establish fixed sequences or give a test, you will also state what user's action is to evoke what presentation.

There is reason to suppose that hypertext systems, once created, will serve a range of purposes. These will include the informing of professionals, education, briefings, giving directions, and even providing leisure reading. In the present SAMI context, this generality is not accidental but necessary. While we will not discuss hypertexts for education in this article, we will discuss the need for this generality in satisfying the customary purposes of information retrieval.

Of course it is risky to speculate on what the particular forms and uses of hypertext will be, since its actual presence and demonstration should swiftly outstrip speculation; until then we can have little feel for it. For any but conventional text, we have little notion of the comparative sense or utility of different sequences and structures. All that we can tell at this time is that there are many possibilities. Until the money is spent to try it, speculation is all we can do. But let us do that.

GENERAL TEXT STRUCTURES

Let us consider in the abstract some of the ways that hypertext may be organized. We will assume the use of some machine capable of presenting the texts in any arrangement, sequence or response network that we wish to plan. What ways are there? (To save space and accentuate continuity with the present, we will refer only to "text structures," recognizing that arrangements which depart from the printable collectively deserve the term "hypertext.")

A. TEXT UNITS

Text structures are composed of units. These units may be for convenience(the page of typescript) or have some meaning (the paragraph; the book.) The units may be composed of other units (as chapters are made out of paragraphs) and may be combined to form larger units (as chapters make up a book).

It is important to note that we may discard conventional document boundaries without discarding the notion of unit organization; organized units of text need not have linear sequence, as they do usually. Moreover, organizing text into units need not restrict the types of interconnection among these units that may be tried. We may indeed define unit boundaries in hypertexts, to give creators and users a sense of orientation, and to facilitate administration.

B. TEXT TYPES

Let us then consider different types of texts. Acknowledging that this taxonomy is quite tentative and uncertain, I will try to enumerate some characteristics that texts of all kinds can have. The object is to learn what traits will distinguish the different kinds of text (especially, hypertext) from one another.

Collection vs. Composition. We may distinguish text structures according to the way in which their units are created. We may distinguish between *collected* texts, which are assembled by some method or rule, and *composed* texts, which are put together with consideration of some of their overall properties by some human agency. As examples of collected text we may take a full mail box, a library, or an issue of *The New York Times;* as examples of composed texts we may consider most books or poems, or the front page of *The New York Times.*

Sequentiality. The next obvious distinction is that of sequentiality. The texts may be *sequential,* that is, have a specific order of presentation, or

structural, with the reader free to move among the parts as he sees fit. (Books as we know them are of both kinds, and there are mixed cases as well; again, the newspaper.)

Connective Structure. The connective relations of texts are yet another basis for distinguishing among them. There may be many different patterns of discrete connections among minimal text segments. Similarly, between the larger text units—which may contain many pieces—there may be many degrees of overlap, interconnectedness and penetration. Print technology has limited our experience of text interconnection, but there are simple conventional examples. These include parallel texts (commentaries, parodies) and hooked-on materials (the footnote, the illustration).

Hierarchical relations may also exist among text structures. That is, some units may be "over" others. By convention this hierarchical relation has to do with conceptual generality in the higher unit. However, hierarchy might take on other meanings. Even for the existing "generality" convention, lattice or other connective structures are possible, where the "more general" units could be for different contexts and settings.

The hypertext will have other attributes, based neither on units nor connectedness. These we may deem (depending on various things) to be "at the user interface," at the option of the text's creator, or left to the decisions inherent in an administrative setup. They include the user's freedom to read or use anything contained in the text complex; his ability to take notes, trace his wanderings, or otherwise to couple the use of the hypertext with his own library system; and permission to contribute to the text.

The supporting facility that presents the hypertext must also have—besides displays and calls, indexing, and the ability to cross-follow from one part to another—such other functions as word searches and counts. (However, note that the provision of service functions that speed up lookups and cross-referencing should not be confused with the hypertext medium.) It may also usefully include such complex indexing and queriable information structures as are now associated with DR and CR, just as books now contain indexes.

SUPPLANTING IR WITH HYPERTEXT

The problem is to answer questions and update specialists; that is, to facilitate specialist access and muster of information (SAMI).

Let us consider here a particular text structure for SAMI purposes, the specialist (or scholarship) hypertext. Let me offer speculations about what the specialist hypertext will be like.

The hypertext for these purposes will probably be a particular kind of

anthology. It will be a repository hypertext, continually enlarged and intermittently edited. Its constituent materials will be from time to time assembled into composed units; for instance, a collection of source texts in a given area of discourse. This collection will be connected with many links to specially-composed review articles, arranged in various strata of summary, detail and commentary. By reading in one of these master texts, and turning to original materials or other instructional sections, the user may keep abreast of a field or instruct himself in an unfamiliar area.

The creation and maintenance of these specialist hypertexts will be a cyclical process. First the "original" units or documents will be put in; these will correspond to reports and articles as they are now conventionally published. So assembled, they will constitute a collected hypertext, structured only by minor indexing—cross-referenced, say, by chronology, authorship and topic. These source materials will remain in the system.

From time to time these source materials will be assembled into composed units; without, however, losing anything in structure or independence, since they will continue to be separately addressable. The new composed units will be created by author-editors, like anthologies and review articles. These author-editors will write overviews, summaries and other passages of their own, create connections, make excerpts and digest versions, and add clarifications, arguments, comparisons, diagrams and miscellany.

The units will be signed, like review articles, and thus linked to the professional prestige of their authors. Short squibs and even connective relations may be attributable in this way to particular persons. The authorship will consist of both writing new material and making connections and arrangements among parts of what existed before. It follows that readers who know something may *contribute* materials. This requires an appropriate administrative context. Some people might need to have their work reviewed, others not.

Such a corpus will have new research results, discussion, and interjections continually added to it by identified individuals. It will be continually enlarged, restructured, reconstructed. Periodically the overview-articles will be updated or replaced, with the old ones kept available for the record. Materials of fading interest will not be destroyed; it will just take longer to get to them, and the newcomer will hear of them less.

Dissents may be incorporated. The system imposes no assumption of completeness or infallibility on its constituent parts. Inconsistencies will be welcome; and with all sides of any argument immediately available, the totality of interrelations among divergent viewpoints will be sooner understood.

Such an approach reflects the real structure of scientific and academic

fields. The text network of current interest will grow and change, rather than, in its fixity, depreciating in relevance as a remembrance of the way thinking went at a particular time.

REMARKS

Curiously, the common view of the information system appears to take it for granted that mechanization will mean depersonalization. "Naturally, such a machine will not be popular among authors since by and large no reference will be made to the originator of any given idea. But, on the other hand, perhaps this will be all to the good since it will avoid the hard feelings which are so often produced at the present time by inadequate referencing."[7] There is no need for such dehumanizing; it will work better the other way. The hypertext approach ties in with existing human institutions and motivations. We must discard the idea that efficient mechanization means arbitrary truncation.

It may be complained that we should eliminate the vagaries and biases that individuals composing hypertexts will introduce. To this I would reply that vagaries and biases which are signed by individuals, and for which they are responsible, are far preferable to those that are silently designed right into the system's heart.

To be useful, the hypertext medium requires some degree of prose discursiveness in the material, or variety in the ways and sequences that the same material can be connected together or explained. Indeed, these texts may be made big and diverse enough for study by specialist and beginner alike, with many entrances, tracks and specially-oriented meanders. Thus the user's previous background and level of knowledge could be taken beneficially into account by the author-editors.

It might be claimed that facilities of such a general character are not necessary for the purposes conventionally served by information retrieval. The purpose of retrieval systems is to inform the technical user of what he needs to know, giving short replies to discrete specific questions.

This may sometimes be the case. When technical queries seek only to fill the blanks in some frozen context, DR or CR will do. Often the user has all the background he needs, and no long explanation is needed as an explanatory frame for the short reply. But even the specialist cannot always ask specific questions or understand short answers, let alone cryptic titles and summaries of borderline literature. For such needs these systems really ought to be ready to teach in the larger sense, even when talking to professionals. SAMI systems need the backup capacity for general instruction or review. And so the function of the hypertext, even in the SAMI context,

will be somewhere between "teaching" and "giving information." (How often there seems no feasible practical distinction between these two.)

Such systems will not, of course, replace all information retrieval systems. Where a fixed-context retrieval structure safely tells the user all he needs to know, as does a train schedule, nothing further is required.

UNDERSTANDING THE MEDIUM

Computer branching operations, text storage and display determine this medium. If this is the case, why is the medium taking so long to be developed? I claim that our problem has been not knowing, in some important sense, what the things are *for*, even though these devices have been with us for some fifteen years.

If complex text display and handling is (or determines) a medium, why doesn't everybody know it?

There are two answers to this, which complement each other.

The first is that various people *do* know it, *have* grasped the medium's character. The second is that there is no reason to expect everybody to catch on to a new medium right away. Let me discuss both of these assertions.

To predict that we will have non-linear text media is not very original. Similar ideas have been expressed by many people, beginning with Vannevar Bush,[8] if we do not mention the Oz books and other immemorial sources. Marshall McLuhan has said it in a very general way: this "mosaic" medium is another step in the transition to organic thinking, learning and society that he predicts.[9] The suggestion here encompasses views expressed by Harold Wooster,[10] H. Bohnert and M. Kochen,[11] T. M. Williams,[12] and (in different places) by several contributors to the Intrex Planning Conference.[13] Bush's prediction was the most thorough. But someone identified only as "Our Special Correspondent" also made the point rather precisely in the *Times Literary Supplement*, May 4, 1962, when he (?) said the scholar of the future would "do more than simply manage linear prose. . . ." Rather, he would "cut it like a film or compose it like a picture. . . ."[14]

CULTURAL LAG

Sociologically, we may regard information retrieval (DR,CR) as a somewhat conservative response by the professional community to various technological innovations. It is a time-honored sociological theory (advocated chiefly by W. F. Ogburn[15]) that practices take time to catch up to inventions. This is called "cultural lag," the delay between a change in society (especially a technological change) and the culture's catching up by squar-

ing related practices with the change. Confusion and inefficiency follow the change until the society finds appropriate ways to deal with it.

For reasons both personal and corporate, the irruption of technology in information handling could not be ignored. It had to be assimilated. "Information retrieval" continued the things that had been done before, extending the indexing and abstraction processes. Document retrieval is an advanced card file; queriable information networks are advanced reference charts.

THE MOTION PICTURE ANALOGY

I would like to draw an analogy between hypertext and the motion picture. By considering the time and trouble it recently took to understand the nature of the motion picture—another prepared temporal and visual medium—we may see better why things are difficult now.

Motion pictures are of interest to us in several contexts; in the development of their elements and techniques; in the development of organized units, especially units which are effective and viable; and in the development of the signatory property, the ability to be signed. These comparisons are suggestive analogies rather than close parallels. Let me recite some highlights from the history of motion pictures, so that the sorts of development that occurred, and what was on the minds of its developers, can enlighten our present. (This will concentrate on the photoplay; let it be understood that other uses of the film are derivative.)

The technology of the motion picture was not suddenly created in the shop of any one inventor. However, Edison's invention, first tried with film in 1889, was the one that took. He called it the kinematograph. It permitted fast viewing of successive frames on a long strip of film. Edison did not work on a projector because he and his advisors thought it would destroy the novelty. He also failed to patent his machine abroad. In 1895 Thomas Armat made a working projector for Edison's film. About the same time, the brothers Lumière, photographic manufacturers in France, created an all-purpose movie machine: the Cinématographe, which could be used by turns as camera, printer and projector. They promptly took it on the road as a novelty, showing films of local scenes, the arrival of trains, and the marching of troops.

The next year, 1896, a real showman began making movies in France. His name was Georges Méliès. A magician and theater operator, he at once saw the potentiality of the medium for presenting the astonishing and fantastic. He made many science-fiction and magic shorts. Méliès originated fades, dissolves, animation, and the whole range of trick photography and multiple exposures. He is also credited with the invention of editing. With *Cinderella* in 1900 Méliès is said to have produced the first story film.[10]

Meanwhile, films in America were already popular, having been introduced to vaudeville in 1896. Credit for being the first "theatrical film" is usually given to Edwin S. Porter's *The Great Train Robbery* (1903). However, Porter's *The Life of an American Fireman,* produced the year before, had contained all the same elements.

A curious monopoly circumstance, with interesting consequences, gradually came about in the American motion picture industry. Because various companies had tried to circumvent the Edison movie patents, there were many lawsuits over the process around the turn of the century. The legal tangle was cleared up by embracing the contending film companies into a monopoly. This monopoly, the General Film Trust, controlled all film production and distribution in America under the umbrella of the Edison patents. For a brief time (1909 to 1915) this trust produced films like sausages, undifferentiated, with the names of the makers and actors not revealed. The officers of this trust preferred that the works they produced be marketable as an interchangeable commodity of a single length (one reel) and quality (mediocre).

D. W. Griffith, the founder of the motion picture art and perhaps the greatest film director of all time, began directing anonymously during this period. In his early work with the Biograph Company, beginning in 1908, he experimented in every possible way with the construction of scenes and narratives, and soon surpassed every other director. Between 1908 and 1915 he invented and consolidated modern camera technique, modern editing technique, and modern narrative technique. By the end of this period his name was known. It was he that consciously began constructing the film out of *shots* rather than "scenes." Griffith's *The Birth of a Nation,* released in 1915, was a "total" film the way none had ever been. Uniting all the techniques that had come before, all his new effects of camera placement, cross-cutting, and tempo of movement and editing, he wove history, battles and love story into a single, entire fabric. Audiences were stupefied. Woodrow Wilson and Lenin were both deeply moved by the film.

Thus the twenty years from 1895 (Armat and Lumière machines) to 1915 (*The Birth of a Nation*) saw the development of a complete medium, in its scope, techniques, and units. No one caught onto it overnight. Nobody saw at the outset what it was going to become.

Now let us examine selected aspects of this brief history. The correspondences to the present situation may be evident.

A. THE MEDIUM'S MISUNDERSTOOD CHARACTER

If we consider the way some of the early things were done, it is by no means obvious that people working in the film at that time had the faintest idea of what was to come. Indeed, it seems that at least some of the film's originators did not know what they were doing or had done. Edison did not want to develop a projection system for his Kinematograph because he thought the novelty would wear off. And when Méliès asked to buy the Cinématographe, Auguste Lumière told him: "It can be exploited for a certain time as a scientific curiosity, but, apart from that, it has no commercial future whatsoever."[17]

Through a misunderstanding of the medium, and a failure to recognize its character, film-makers repeatedly imported to it the techniques of the stage. While stage *content*, when adapted, has worked, stage techniques have not worked in films, and have given satisfaction only for short periods of time.

The use of flat backdrops, action in a plane facing the camera, and the continuous take throughout a whole scene were all early conventions abstracted from the stage. The dramatic scene of early films was called the "action tableau," but it had relatively little action; its flat arrangement was a simple transposition from the stage, and its slight motion a prudent simplification of ordinary staging. This was the technique of Méliès, who also used painted backdrops shamelessly. Though the action tableaux, after Porter, became infused with more and more action, heavy stage influence remained. Even Porter was not completely won over by his own inventions. Though he had pioneered three-dimensional action and stirring visual effects in *The Great Train Robbery*, when he filmed *Uncle Tom's Cabin* he reverted to tableaux, flat sets and flat action.[18]

Stage plays have often been filmed as is; this occasionally happens even in the present day. It does not work very well. In 1907 the Film d'Art Company, in France, was established to uplift the motion picture by filming plays. Their product, while hailed for being cultured, was not very good. In 1912 Adolph Zukor's *Queen Elizabeth* opened. It was an attempt to immortalize Sarah Bernhardt on film. It made money. It also showed rather clearly that the sorts of posture, makeup, movement and declamatory style thought appropriate for the stage were not at home in the new medium.

But Griffith clearly saw the differences between stage and film. He was asked whether he thought knowledge of the stage was necessary to direct films. He said,

"No, I do not . . . The stage is a development of centuries, based on certain fixed conditions and within prescribed limits. It is needless to point out what these are. The moving picture, although a growth of only a few years, is boundless in its scope and endless in its possibilities. . . . The conditions equally dissimilar. . . ."[19]

B. TECHNIQUES

There is a basic roster of techniques of film. These are effects (stop-motion, animation), display arrangements (camera angles and movement), and the selection and interconnection of parts (direction, editing).

These techniques of film are basic and perfectly known to those who make films, though there is disagreement on style and appropriateness. All the techniques are familiar. Nothing is new, in the sense that it is impossible today to create new elements. There has been virtually nothing new in the last 50 (fifty!) years, since, say, Griffith's *Intolerance* (1916).* But these techniques of film were by no means obvious. It took twenty years of film-making to discover them, from 1895 to 1915. Though many of the films have been lost, in principle each innovation has associated with it a specific film, director and date. The various techniques were introduced one by one, chiefly by Méliès, Porter and Griffith.

Early movie-makers did not use more than one shot in a scene. They did not move the camera. They instructed their actors to wave their hands around and exaggerate every gesture, since the camera was so far away. And the camera was so far away because film-makers thought it looked wrong not to show the actor's whole body. Other rules included:

"1. Every scene must begin with an entrance and end with an exit, just as on the stage.

"2. Players must face the camera and move horizontally, except when the movement was rapid, as in a chase, or prolonged, as in a fight. At these times action was in diagonal relation to the camera in order to give the players more area.

"3. Any action in the background must be slow and greatly exaggerated so that it would 'register' on the audience.

"4. Pantomime must be exaggerated and over-deliberate; e.g., a stare had to be held; a start must be violent; speeches must be mouthed with pronounced slowness."[20]

C. UNITS AND THEIR ORGANIZATION

The character of basic units, in particular the narrative film, was not at

*Sound, color, etc., did not modify these techniques; but they added separate problems of their own and additional opportunities for synthesis.

first grasped. "In 1896, a special outdoor performance by Joseph Jefferson of *Rip Van Winkle* amidst the natural scenery of his summer home on Buzzard's Bay, Massachusetts, was recorded by Biograph on 219 feet of film. But instead of shooting this subject with a view toward an interesting and a dramatic continuity, it was filmed in eight episodes billed and sold as separate catalogue numbers. . . ."[21]

For a while, projected films were timed to the vaudeville turn, and then to the organization and taste of the nickelodeon theater. Though films tended naturally to become longer, after about 1908 The General Film Trust held the length of all Trust films to one reel (1000 feet, about sixteen minutes). It was not until the decline of the Trust that the American film could break out of this one-reel mold and attain feature length.

Not only was the overall unit misunderstood; the parts themselves were not thought to be manipulable and discrete. The basic unit was thought to be the "scene." It was only D. W. Griffith who learned to work in terms of *shots* rather than complete scenes; and to construct the scenes from the ground up, using the shot as the building block. Griffith discarded the principle that a scene must end before a new shot may begin, and, indeed, the principle that one scene must end before another begin.[22] By manipulating the constituent parts he learned to handle the motion picture as a whole, controlling and integrating all its aspects.

The climax, of course, was Griffith's *The Birth of a Nation*. Though panoramic spectacle with three-dimensional sets had been used before, especially in the Italian *Quo Vadis*, it had not previously been integrated with a well-told story. *The Birth of a Nation* was such a complete, unified, spectacular achievement: the true feature film, and a model seldom equalled.

D. THE SIGNATORY PRINCIPLE

"Films are made by men, not machines or corporations."[23] Yet interestingly enough, the notion that films should be *signed*—that is, have personal credit go to their creators—was not acknowledged for a time. Quite the contrary: credit was for a time *suppressed*.

The film trust's management of the motion picture industry from 1909 to 1915 produced an even stream of one-reel films whose actors and directors were not known to the public. The public was developing favorites, and demanding their names—like that of the girl everyone knew as "Little Mary." The exhibitors, too, wanted the names so they could in some measure control public response, i.e., book the favorites and advertise. But the favorites' names were kept a secret. (The film trust's policy of not releasing the names of actors was ably met abroad by the simple expedient of assigning names to them.[24])

The cat got out of the bag in 1910. Carl Laemmle of the trust-defying Imp Company publicly advertised the name of a film actress for the first time, after he had stolen her from a trust studio. Thus began the star system. Competition from Laemmle and other outlaw producers forced the trust companies to reveal names. ("Little Mary" found it expedient to adopt the name of Mary Pickford.) Griffith's name was first advertised when Biograph gave in to the new open-names policy in 1913.[25]

IMMANENCE

I have attempted to point out some curious similarites between a general-purpose medium we all know and a general-purpose medium that none of us, to my knowledge, has seen.

The comparison would be strongest, of course, if we based it on a premise of structural determinism. This would be to assert that media like the hypertext and film derive their character from intrinsic principles and internal dynamics that are somehow built right into them.

This was the view held by Kuleshov, the great Russian teacher of film. He said: "In every art there must be first a material, and secondly, a method of composing this material specially adapted to this art."[26] As the reader sees, it took several decades to discover these for the film.

Perhaps in general we have no leeway, then, and media are ours merely to discover rather than to invent. But I am not sure that for the hypertext it is this tight. We can in hypertexts establish precedents and traditions, doctrines and schools; we can, by stress upon programmed instruction, flash card effects or civil liberties, influence for worse or better its usefulness to our society and condition.

The medium has an internal dynamic, I think, immanent and intrinsic in the technology. This we must discover. But however compelling this internal dynamic may be, there will be play for some adjustments. We must strive to add the touches that suit it best to the way we want to do things.

A DISANALOGY: STANDARDIZATION

By luck, movies are standard. The width and sprocketing of Edison's original films became and remained the 35mm motion picture standard. While there are other standards for other purposes (8mm, 16mm) and other rival processes (3-D, introduced in 1922, and such other intermit-

tent specials as wide-screen and smell track), in general, a movie is a movie.*

In order to try out the new medium, other than haphazardly, standardized settings must be created, so that work in this medium can be used and passed around. We need standardized character-sets and connector notation, standard ways of scaling down a presentation to a sparse facility and of interfacing it through varying screens, buffers and accessories. This might be accomplished by the specification of standard facility levels, so the creators of such things could either stipulate a minimum system required for use of a given text, or an appropriate presentation level for each configuration of hardware.

Moreover, it will be necessary to establish reliable royalty provisions for authors, editors and contributors, probably based on the presentation frame. But something equivalent to *ownership of a copy* will also be necessary, so that a given user may have unlimited access to his own possessed material.

CONCLUSION

Information retrieval has extended traditional techniques to new technology. An analogy has been found in the history of motion pictures. The prior techniques of the theater were at first extended to movies, and it was several decades before we understood the motion picture as a medium. It was as natural to extrapolate these prior techniques to the movie as it now seems natural to extrapolate prior techniques of information handling to the CRT-computer-mass memory system. Here again, as in films, the technology makes possible a medium which is radically new and has dynamics of its own. The most suitable techniques and units were not immediately evident in movies, either.

The conventional "document" is not God-given, and in fact is inappropriate for most purposes. Systems based on discrete and isolated documents relinquish the greatest power of the new technology. But at the other extreme we are wrong to suppose that an information machine can or should eliminate the human task of composition. We have not begun to explore the possibilities of natural language woven into more complex (but also natural) arrangements.

The problem of Getting It Out of Our System, then, is not the problem

*It is important to note that facilities are readily available for the transposition of off-standard work to the main standard, by optical printing and sound transfer. Often novelty effects offered in first-run theaters are scaled down to the standard for distribution.

of fishhook design for a document pool, or of creating a conversational black box with a narrow vocabulary. We must get out of our system the fixities of thinking and procedure that hold us back.

CITATIONS

1. J. C. R. Licklider, *Libraries of the Future*. M.I.T. Press, 1965.
2. R. F. Simmons, "Answering English Questions by Computer," *Communications of the ACM* 8:1 (January 1965), 53-69.
3. Thomas Marill, "Libraries and Question-Answering Systems." Report Number 1071, Bolt Beranek and Newman, Inc.
4. Bert F. Green, Alice Wolf, Carol Chomsky & Kenneth Laughery, "Baseball: an Automatic Question-Answerer" in E. A. Feigenbaum and J. Feldman (eds.), *Computers and Thought*. McGraw-Hill, 1963, 207-216.
5. Theodor H. Nelson, "A File Structure for the Complex, the Changing and the Indeterminate." *Association for Computing Machinery Proceedings of the 20th National Conference*, 1965, pp. 84-100.
6. Theodor H. Nelson, "The Hypertext." Paper presented at 1965 Congress of the International Federation for Documentation, Washington, D.C., October 1965.
7. Andrew W. Booth, "The Kinds of Machine Now in Use" (No. 4 in the series "Freeing the Mind"), *The Times Literary Supplement*, April 13, 1962, p. 252.
8. Vannevar Bush, "As We May Think." *Atlantic Monthly* 176:1 (July, 1945), 101-108.
9. Marshall McLuhan, *Understanding Media: the Extensions of Man*. McGraw-Hill, 1965.
10. Harold Wooster, "A Web of Words." Keynote Address of Sixth Annual Institute in Technical and Industrial Communications, Colorado State University, Fort Collins, Colorado, July 1963.
11. H. Bohnert and M. Kochen, "The Automated Multilevel Encyclopedia as a New Mode of Scientific Communication," in Manfred Kochen, *Some Problems in Information Science*. Scarecrow Press, 1965, 156-160.
12. T. M. Williams, "Some Logical Aspects of Thematic Representation with Explicit Ties." Paper presented at 1965 Congress of the International Federation for Documentation, Washington, D.C., October 1965.
13. Carl F. J. Overhage and R. Joyce Harmon, *Intrex: Report of a Planning Conference on Information Transfer Experiments*. M.I.T. Press,

1965. Two notes by J. C. R. Licklider are of special interest: "An On-Line Information Network," 147-155, and "Proposed Experiments in Browsing," 187-197. See also E. S. Proskauer, "The Motivations of Authors—Intellectual Property and the Computer," 199-201.

14. Our Special Correspondent, "Poetry, Prose and the Machine" (No. 7 in the series "Freeing the Mind"), *The Times Literary Supplement*, May 4, 1962, p. 310.
15. William F. Ogburn, "Cultural Lag as Theory." *Sociology and Social Research* 41, 167-174. Jan-Feb 1957.
16. Richard Griffith and Arthur Mayer, *The Movies*, Bonanza, 1957, p. 6.
17. Maurice Bardèche and Robert Brasillach, *The History of Motion Pictures* (tr., ed. Iris Barry), W. W. Norton Co. and the Museum of Modern Art, 1938, p. 10.
18. A. Nicholas Vardac, *Stage to Screen: Theatrical Method from Garrick to Griffith*. Harvard University Press, 1949, p. 171.
19. Lewis Jacobs, *The Rise of the American Film: a Critical History*. Harcourt, Brace, 1939, pp. 118-119.
20. Jacobs, *op. cit.*, p. 60.
21. Vardac, *op. cit.*, p. 169.
22. Ernest Lindgren, *The Art of the Film*. London: Allen and Unwin, 1949, p. 71.
23. Iris Barry, quoted by Richard Griffith in "Moving Pictures," *Encyclopedia Americana*, 1960 edition, v. 19, p. 540.
24. Griffith and Mayer, *op. cit.*, p. 47.
25. Jacobs, *op. cit.*, p. 88.
26. This is Pudovkin's reminiscence, cited by Lindgren, *op. cit.*, p. 76.

(64)

(reprinted from:)

COMPUTERS IN INSTRUCTION
Their Future for Higher Education

PROCEEDINGS OF A CONFERENCE
HELD IN OCTOBER 1970
SPONSORED BY THE
NATIONAL SCIENCE FOUNDATION,
THE CARNEGIE COMMISSION
ON HIGHER EDUCATION,
AND THE RAND CORPORATION

R. E. LEVIEN, EDITOR

R-710-NSF/CCOM/RC
JULY 1971

Rand
SANTA MONICA, CA. 90406

(65)

COMPUTOPIA AND CYBERCRUD

Theodor H. Nelson
The Nelson Organization, Inc., New York, New York

Here I am again. I regularly cry in the wilderness at banquets, conferences, and other ceremonial events. I have been doing my thing for ten years, and giving these talks and publishing articles for five, with little result. My articles have apparently been boycotted by *Computing Reviews*. I seem to be the Electronic Eeyore.*

The truth is that I am not really trying to interest computer people. They are the ones who only seem to tell me, "It can't be done." There turns out to be no answer to that. Many times I have come up with a way to do something, in which case the man who said it couldn't be done usually says, "Why would you want to do such a thing, anyway?"

It is the laymen and literati, the noncomputer people, I want to reach. But this also seems futile. A few laymen and literati occasionally seem to become quite turned on; but without Experts to confirm what I say is possible, they gradually edge away and don't listen either. I feel like Marco Polo in his later years, no longer Italian and certainly not Chinese, trying to interest one in the other.

I am here again to say what I always say, not so much to influence anybody now as to say "I told you so" later. I would rather say what I think, and seem to be a lunatic, than try to make a favorable impression by lying about my true opinions. This way also I can treat you to a shock of recognition several years hence, when you say to yourself, "Oops—that was the prediction."

In 1966 I spoke at Rand, also under Roger Levien's sponsorship. The 1966 talk, "Hard and Fast Thoughts for a Softcopy World," is perhaps interesting in several respects. I made a number of crazy predictions. They seem less crazy now, and various of them have come true already. I have nothing to retract and relatively little new to say, except for a few details.

*See especially "A File Structure for the Complex, the Changing, and the Indeterminate," *Proc. ACM 1965;* and "Getting It Out of Our System," in Schechter (ed.), *Information Retrieval: A Critical Review*, Thompson, 1967. The tale is poignantly told in "Barnum-Tronics," *Swarthmore College Bulletin* (Alumni Issue), December, 1970.

(66)

What Roger has asked me to talk about again tonight is "the computer as a medium." But that doesn't quite put it my way. I see the computer as the heart of the new presentational systems of the future—and this vision seems curiously different from what we generally hear. What we must work on are the principles of presentation, not of computers or other technology.

We are usually told, on various sides, that some kind of revolution of human information is upon us, but somehow in the course of things the computer will make this revolution inhuman. Chugga-chugga, we will have to learn obtuse query languages. Rattle-clank, we will get answers back in symbolic logic. (Too bad for some of us.) Tippity-tap, the terminal will tell us a thing and then ask us what it told us. Tell and test, tell and test, instant boredom, but who dares argue with science?

The title of my talk is a slanted dichotomy: "Computopia and Cybercrud." However, these terms may need a little elucidation.

"Computopia" is of course a portmanteau of "computer utopia." Each of us in computers has (or should have) his own computer utopia, some vision of a future better world in which the computer figures prominently.

Unfortunately, there are many trivial horizons being offered the world by computer people: shallow possibilities, uninformed aspirations. I refer not only to commercial enthusiasms—one sees published articles with titles like "New Horizons in On-Line Credit Reporting"—but also to cramped visions of the ideal. Cramped visions are all too common; but the worst thing is that here they are contagious. They are sold to the public as technically necessary, and the public doesn't know any better. This is the cybercrud problem: advice and creation of systems, supposedly based on technical requirements, whose categories and rigidities are unnecessary. In the worst cases they are not only unnecessary but wrong.

If the public doesn't know what is technically possible, you can sell them any inflexible kludge and tell them it has to be that way. Unfortunately, technically sophisticated people forget how uninformed and timorous about such matters the rest of the world is.

Once I was surveying individual departments in a big company, having been assigned to find out if any new technologies could be used in the work at hand. Now, I couldn't just ask, "How could technology help you if you knew more about it?," since that's like asking, "What fact that you don't know is the most important?" So I developed a little song-and-dance for the purpose. "Suppose," I said, "that anything was technologically possible. Suppose it's the year 2000." (This premise was the tongue-loosener. Somehow the year was sufficiently unimaginable, or the sense of having to be answerable for the consequences could be suspended.) "Suppose technology could give you anything you asked for," I said. "What could your department use?"

The interesting thing I learned was that when you asked mature people in business what they might want in the year 2000, they almost invariably (a) wanted computers to take dictation, but otherwise (b) came up with things which were presently possible, though perhaps expensive or awkward. But one reply I will never forget came from a lady in charge of a language teaching program. "I suppose that by the year 2000," she mused, "each child could have his own tape recorder."

The horizons I am talking about lie a little further. Within the coming decade we will see the explosive growth of computer display, an expansion that will rival

or surpass that of television, and compare in ubiquity to the very telephone. We are going to have an entire cultural revolution based on computer display.

It is my belief that many important benefits can flow from this revolution, if we do it right. I believe enlightenment, knowledge, and understanding can be furthered throughout the public. I believe creativity can be fostered in many of its forms. And I believe a new and important freedom of information is possible.

This is my prediction and my call to battle; evidently few of the most ardent enthusiasts of computer display go so far. But if it is correct, it means a revolution in human life and thought comparable to what followed Gutenberg. My interest is in giving shape to that revolution, in urging it toward enlightenment and humanist freedoms, rather than having it stumble accidentally into the formalization of dreariness.

There is no overall term for what I am talking about. "Information retrieval" and "computer-assisted instruction," with their false ring of exactitude, say less than they seem to. "Computer display" is a technology. But I would like to propose a term for what I am talking about: *fantics*, the art and technology (in that order) of showing things.* The major precepts of such a field should be to make things look good, feel right, and come across clearly. The use of screens and computers is indicated.

What will these screens be for? And what will be on them? Many people seem to be in the grip of the travel-agent idea—that screen terminals will be for bank clerks and airline reservation people. Sure, for a while. But they are also going to be in the home. (And if we have them in the home, they may not be needed by travel agents.)

I want a world where we can read the world's literature from screens rather than personally searching out the physical books. A world without routine paperwork, because all copying operations take place automatically and formalized transactions occur through formalized ceremonies at consoles. A world where we can learn, study, create, and share our creations without having privately to schlepp and physically safeguard them. There is a familiar, all-embracing motto, the jingle we all know from the day school lets out, which I take quite seriously: "No more pencils; no more books; no more teachers' dirty looks." The Fantic Age.

If you asked me what my computopia is, then, I would say a sort of Woodstock with terminals. With terminals on all sides, we can more easily go barefoot and pocketless. I propose to turn on people's minds with display screens rather than drugs.

Those of you who have moved your possessions and papers while switching universities are aware of the burden on the individual of retaining documents and records in paper form. This weight off our shoulders will be a pleasant side benefit. But I see two greater steps we will be taking. Our screen-and-computer systems will help us in the difficulties of organizing thought, of revising writings, of indexing complexities. And when our works are finished, we will send them out on screens —marvelous new forms of writing and illustration.

In dark contrast to such a possible enlightenment, I would like to point out an

*The term, etymologically impeccable, is from the Greek *fanein*, to show. The pleasing connotations of "fantasy," "fancy," and "fantastic" are fringe benefits.

unfortunate tendency, occasionally a villainous practice, which we may call *cybercrud*. By *cybercrud* I mean putting things over on people using computers.*

Cybercrud can take many forms, all related. The computer's cachet may be used to hide your premises, the way you want to do things, the secret loadings of your approach and procedures. The computer, its accessories, and terminology, can give the semblance of validity to all sorts of procedures or statistics. The term "computer" is to many a rubber stamp meaning "scientific."

We may use the computer, or the mention of it, to perplex, intimidate, or bamboozle. This, too, is cybercrud. The following subscription renewal letter has a certain charm.

> Dear Subscriber:
>
> Because [our] . . . subscription roll is maintained by electronic computer, it is necessary for us to assign a common expiration date to all subscriptions. This enables us to distribute copies and mail renewal notices to all subscribers at the same time. Therefore, we are writing to inform you that your . . . subscription must be renewed *now*. . . .

We may note that while this electronic computer requires all subscriptions to end at the same time, it nevertheless does not require all subscribers to live at the same house or have the same name.

Cybercrud is most surreptitious, even if benignly and sincerely, in the "one-way" form. This consists of presenting one's own way of doing things as *the* way, hiding the fact that it is one of many alternatives. Of course it is natural, in the throes of enthusiasm, to forget to mention alternatives, to present as scientific the consequences foisted by one's premises. "This is how the computer bakes a pie!" says the press release—and recipes are not compared.

With the public so flummoxed by computer news, and so easily taken in by breakthrough pronouncements, the mischief of this tendency is great. "You have to submit out of technical necessity" is the gist of this form of cybercrud. This is equally unfortunate when said by computer salesmen or by professionals prematurely set in their ways.

Of course, it is possible to believe sincerely that one's own way of doing things is the computer's way, the only way it can be done by computer. But we know that this is rarely if ever so. Actually the computer is a *tabula rasa*, a projective system, whose behavior takes on the preconceptions, and sometimes the personality, of those who set it up. But the possibilities can far exceed what given individuals think of. And this is the problem that brings us here tonight.

*Since in the course of present-day technology promotion we find ourselves affronted with too many cruddy words beginning with "cyber-," I thought it was time to coin my own "cyber-" word, one that would be explicitly, rather than implicitly, cruddy.

The derivation of "cybercrud" is as weaselly as its meaning. To the general public, the word "cybernetic" means, rather incorrectly, "related to computers," whereas from a technical point of view it more exactly means "concerning control linkages." Thus the prefix "cyber-" has the proper connotation of "spuriously related to computers." The actual *Greek* "cyber-" meant "steersman." Thus at the micro-level we may think of "cybercrud" as meaning, "steering into crud, with spurious connotations of computerishness."

EDUCATION

That school is stupid, boring, and insulting need scarcely be mentioned, except that we tend to forget it. We forget the inanity, the complete nonsensicality of most grade school and high school pursuits. Their insignia of officiality somehow seems to make them right. What matter if children's good time is being wasted by astrology, candle-dipping, or Euclidean derivations? Any nonsense will do, so long as it leaves a trail of *grades*, evaluations etched in history that can be used to blackmail the victims or their parents.

Curricular timing and grading virtually obliterate the nature and natural interest of every subject. It is as though we were taken in groups to visit national parks in the back of a truck, racing on a treadmill. Every possible activity is related to a made-up standard; nothing is allowed to be merely interesting. For those on the up side, the system furnishes clues as to direction of reward; for those on the down side, the presumption of *own failure* is affixed to the subject. Few of us do not learn that we are "no good" at something, and adult regrets are heavy with both the realization that it could have been different and self-blame that it was not.

Some remarkable traditions govern the structuring of subjects as we teach them. But they are no more in the nature of the things to be learned than the cuts of steak are in the anatomy. Every subject has a beginning, a middle, and an end; it is laid out by the assigned reading and precisely bounded by the scope of the final exam; every topic may be reduced to shallow enumerations strung on vague explanatory connections, dismally explained inanities, explanations which associate phrases without a sense of meaning, and incompletely explained "skills" to be practiced without insight. Questions and other matters therefore become either "relevant" or "not relevant," according to whether they fit the boundaries and sequence of the "subject." And thus it can come about that the answer to a question is, "Just learn what you're told."

We have been misled into believing that all this is how it has to be: the cascade of premises flows into a landlocked swamp.

The idea of prerequisites and "basics" is one of the few justifications for having curricula at all. In a few clear cases—say, the differentiation of mathematical functions, or punctuating correctly—it is impossible to do the more complex detail work without having learned exactly to do the simple detail work. But I believe that such cases have grossly misled us. We suppose that because some learning sequences cannot be circumvented, then all learning should be reduced to sequences. This has several disastrous consequences.

First, unique curricular paths. The curricula of the schools are generally designed as pathways radiating from some primal state of ignorance, wagon-wheel spokes without the wheel. The only way to learn most things in school is by taking an exactly prescribed series of intermediate steps. There is no way around. And if one has not taken the steps between, a thing simply cannot be learned.

The psychological consequence is most pernicious. There are things one "knows" (if the details are forgotten, one is still oriented) and things one "does not know" (there has been no introduction, no orientation). A sense of weakness is produced; one drops a subject that has come up; one shrugs. But this is only part of it. What is much worse is that when one has "failed" in a subject, all further thoughts about this subject are darkened, colored by this sense of failure. One avoids;

one strives actively to find other interests or distractions: "Aw, never mind that stuff."

Last, the sociological consequence. We produce people with funneled minds, the so-called "types"—the literary type, the scientific type, the mathematical type. And an occupational structure around these types. And subcultural divisions. And everybody stays where he feels safe, and everything interesting in the world is hidden from almost everybody.

A lot of people I talk to say, "I don't think it's as bad as you say, but I agree that the educational system can be improved." But I'm not sure it *can* be improved. The system is geared and swivelled to do exactly what it does: discourage and mediocrify in the guise of sticking to business. It's not merely a question of whether real improvement is possible within the system, but also of whether it would still be the system if we made any real improvements. Imagine (if you can) people growing up without a self-stereotype of structured disabilities, who thought they could do anything if they worked at it. Imagine most people having a real choice of occupations. Imagine if the classroom atmosphere could be stripped away and the students put in genuine rapt communion with the subject—a state which teachers are sometimes deluded into thinking they achieve—without the deleterious effects of competition, time pressure, and stigmata for real participants. Imagine kids experiencing the excitement of intellectual issues—and not just as a fleeting part of that grim academic exercise, writing a paper. Imagine students actually interested in school. Or whatever we would now have to call it.

DESIGNED AND SIGNED WORKS IN COMPUTER MEDIA

Roger Levien says that the part of my 1966 talk he remembers is the discussion about "the computer as a medium." Let me go into that in some detail.

I said this: A number of new media will come about, employing computers or other digital means. The computers will store, respond, display, and follow complex directives about their response. What they show will consist of *works*, intentionally organized things whose impact on the viewer or participant has been designed, much like the parts of a picture, a book, or a movie.

By "media," then, I mean stabilized forms of presentation that people will create works in, for other people to use and enjoy. The presentation may be carried out by computers; but its plan, and the viewer's feelings, will be mediated by other human beings.

We have many media now: the newspaper, the movie, the phonograph record, the TV show, and many more, each with its own variations. We understand these media and the people who do things in them; we understand the position of the reporter and the columnist, the photo-journalist and the movie director, and their contributions in bringing us facts, impressions, and visions.

Yet for some reason we have failed to extend this understanding to the media that will be brought to us by the computer. There is a floating myth that computer media will be different from all those that have gone before, either thrown together on the fly by the machine itself, or presenting passive nonedited descriptions of the world, or dutifully constructed on scientific principles by psychologists who remain

aloof from the content. I find all of these ideas rather absurd, especially since to formulate them we ignore our extensive experience with other media.

If the computer itself puts together presentations for us, it is acting not as a medium but as a facility, according to rules we will wish to control. (Even if it presents a montage or pastiche, the locus of action is still not the computer, but the formula for what the computer is doing.)

The second view is that the computer and display will merely dip into some data collection which passively describes the world. One new machine, the Evans and Sutherland LDS-1, dips into a three-dimensional data structure, much as a goldfish net is dunked to surround guppies. The user may roam and wheel through whatever world is represented in core memory, with almost no programming needed. Such systems might leave one thinking of the material to be shown as simply a view of the world, a neutral representation or assortment, to be routinely prepared by anonymous coders. (Simply to describe it this way shows the suspiciousness of this view.)

The third possibility is a view that there are neutral or colorless scientific methods for organizing presentations, especially those methods growing out of learning theory. This seems to me like saying that since photography is optical, photographers should be trained in optometry, or that record producers need to study psychoacoustics. It is not necessarily relevant to what should be done or what the media are going to be about. If the principal impact of a teaching system is a matter of feelings, appeal, and cognitive handles, as I believe, then *these* are where we must place the emphasis.

But this is the general view in "computer-assisted instruction." That instruction must be the presentation of sequenced chunks and questions in a conversational format seems to have become a fold premise, generally blotting out other possibilities of computer teaching. At best, computer-assisted instruction can be very good indeed, in which case the only possible criticism is of the immense cost of preparing materials. But at worst, it seems to be a conspiracy to do some of the worst things school has ever done: cut and dry the material, spoon-feed, bore, and insult the participant.

I think none of these models represents the main thing that will happen. Rather, people will be designing display-based media and creating works in these media. Indeed, people will be signing them, just as ever before, and viewers will seek out the works of particular authors much as we do now in every other medium.

We have gotten to my position and main interest. I believe the computer's astonishing possibilies as an aid to the human mind and imagination have scarcely been touched, and the new media made possible by the computer and display will be miraculous and awesome.

I have suggested using the prefix *hyper-* in general for multi-dimensional and nonsequential computer media, those having some kind of complex order within which the user may roam.

Since computers can control any other equipment, we may expect new media which tie together old forms of presentation in new ways; for instance, the branching motion picture (hyperfilm), or branching audio. But I suspect that all these hybrid systems will be comparatively unwieldy—the expense of branchable film transport systems, for instance, is immense compared to all-digital systems—and interest will converge rather soon on the pure digital media for the computer's own display: the hypertext and hypergram.

(72)

PRESTIDIGITATIVE PUBLISHING

To supply our scopes with the hypermedia we will want to read, I foresee a new era of publishing, and a whole new publishing industry. In this coming era the digital files of the publisher will be connected by telephone (or other means) to the subscriber's console. As you ramble through hypertexts or explore hypergrams, the news of your actions will be flashed to the great feeder machines at the publishers' distribution centers. These feeder machines will disgorge to the customer the furtherances of what he is doing, keeping him continuously supplied. The material will be copyrighted, and small royalties will be continuously billed for screen-minutes of presentation. Entirely new creations and the works of the past will be equally and quickly available. Enchanted gardens of information, prearranged by authors and editors, will be available to you. Screen pyrotechnics and display tricks will be intertwined with pictures and text. There will be anthologies, magazines, encyclopedias—or things like them.

A related development will be *creative facilities*, systems for using computer displays to help in creative activity. Publishers will serve as storage warehouses for the overflow from such systems, and probably buffer materials and messages being sent privately among users.

The user will be able to perform input and complex parallel annotation, keep a trail of his own activities, and backtrack through his own past actions and those of others. In other words, Vannevar Bush's "memex" will come about substantially as he saw it."*

STRETCHING THE IMAGINATION

An important human tendency is to rationalize and order what we hear. I have found that people at my lectures continually assimilate what I say about hypermedia to some pulsating "new world of communication" they have heard about, mushing together computer-assisted instruction, information retrieval, holography, and even cartridge TV.

I often talk about stretchtext, not because I am so attached to it, but because listeners who understand the concept cannot continue to believe I am talking about the usual stuff.

"Stretchtext" is a form of hypertext I have suggested for discursive written subjects, such as history and the social sciences. In stretchtext, the reader may control the amount of detail to suit himself; as he pulls on a throttle or some other control, additional words and phrases appear on the screen, and the rest move apart to make way; as he pushes the throttle in the other direction, words and phrases disappear, and the rest of the text slides back together.

The presentation should never change sharply. Smooth motion of the text pieces on the screen is utterly essential. For instance, a long sentence should, as expansion continues, be broken into two short sentences as similar to their parent as possible.

*The particularities of this prediction are pursued in my article "As We Will Think," presented at the September 1968 meeting of The American Chemical Society; to be published.

(73)

Although we have as yet no firm evidence, my slight experience trying to write stretchtext suggests that it is no harder to write, and perhaps easier, than ordinary prose. To each small piece of text is assigned an *altitude*, or "degree of stretch," at which it is to come and go. The problem is basically to lay out an overall expository structure and then find some way to compress it gradually without large jumps or reversals of sequence.

Discrete hypertexts and hypergrams are essential and basic. Simply put, these are writings and pictures with footnotes that extend in all directions. Touch an asterisk, or perform some equivalent act, and the system will bring another connected thing to the screen.

More complex hypergrams are very desirable, particularly those capable of immediate and smooth response on the screen. Consider, for example, a map on your screen which you may move around the world and zoom in on any part, to any degree. Look at the whole United States and zoom in on Secaucus, New Jersey, and environs. By various manipulations the user should be able to overlay this picture with populations, climate, historical events, light industry, or crime statistics.

Consider another hypergram: a picture of the human brain which the user may actually manipulate in three dimensions. (This is a particularly interesting example because of the stereoscopic complexity of the brain.) The user should be able to rotate, magnify, zoom in, change perspective and viewing angle, brighten different subsystems, and obtain annotations.

Finally, consider a cartoon human body on the screen, having various simple animations (heartbeat, peristalsis, etc.) and permitting the user to open it up in various ways and views, as well as get labels, explanations, and other annotations.

It will be noted that in none of these media have I made any reference to sequence. Thus it will be plain that these activities are not the same as the prescribed-sequence formal activities being pursued generally under the name of "computer-assisted instruction."

Finally, it may be divined that these systems overlap in function with both "information retrieval" and "computer-assisted instruction." I am persuaded that neither discipline, in its larger sense, can or should have a separate existence. "Information retrieval," unless it deals with the simplest facts, must also involve orientation and learning materials. And, as I have already pointed out, systems intended for teaching will need to open up whole new areas of option, until they become very, very like hypertext.

INTENTIONS AND PRIORITIES

The next question in your minds may be, "Where does all this lead?" If instructional sequences are not the right place to put our effort, what sort of things do we do next?

There are two answers. On the one hand, we should be trying out new organizations of information, grand complexes of words and pictures that the user can peruse, explore, and learn from. Unfortunately, present presentational systems are very badly suited to this purpose. Conventional computer systems are designed with other things in mind, and I am increasingly convinced that some rather basic changes are needed before proper environments are available.

(74)

new approach, utilizing through-designed hardware and system programs. I will shortly talk about some of the approaches we are looking at.

DESIDERATA

Keyscopes, or text display consoles, are not yet what they should be. While the more attractive units, such as the Datapoint 3300, are now comfortable and attractive enough for long hours of use, they are not yet sufficiently zingy for the hypertexts we are going to want. The most important feature such text consoles will need, in my view, is some capability for smooth text motion. Only if you can see where a piece of text is coming from, and which way it goes, can your eye and mind remain oriented. (Unless the display space is trivial, say, one made up of "pages," or long scrolls.) Moreover, variable-sized characters, italics, other fonts, and serifs are all desirable, in that order.

It is possible that keyscope terminals, such as the TV raster-chopping type, might be improved along these lines, but the engineering might be difficult (and the engineers unmotivated). For these reasons I put my hope in calligraphic systems—those which draw with a programmable beam of electrons, such as DEC's 338, IBM's 2250, the Adage terminal, the Vector General display, etc.

And for the general purposes of CRT hypermedia, even the best, excepting always the miraculous Evans and Sutherland machine, is not yet good enough. The finest calligraphic displays, such as the 225° and Adage terminals, are usually hampered by the big computers that support them. A sudden refusal to respond often means the main computer is stopping to print something or punch somebody's cards.

Some way has to be found out of this. The usual solution proposed is to have more and more immense computers doing the main work for the display. The solution I favor is to use the display's little computer better and use big computers less or not at all.

XANADU

The term "Xanadu" I have used since 1966 as an ongoing project name for the console I have wanted to build. The overall set of ideas and preferences has been churning in my head since 1960. Its functional specifications, frozen along with its name, now seem within reach. I called it "Xanadu" after the pleasure dome in the poem "Kubla Khan," with its connotations of mysticism and artistic trance, to say nothing of the cachet added by Orson Welles' appropriation of it as the name of Citizen Kane's palace.

Here was the idea of it:

1. It was to be created for the naive, antinumerical, and accident-prone user. There were to be no visible numbers, save those explicitly put in by the user. Unintended actions by the user could not harm his files or undo his work, because of several levels of backup, fail-safe, and warning. (At the extreme, procedures of the utmost gravity would be required for dangerous and irreversible actions, similar to the protocol said to be established for the firing of intercontinental missiles. For

(75)

COLLATERAL DATA STRUCTURES

First I recognized that a certain relationship was essential and basic. I have called it different things at different times, "zipper list" being the simplest. The idea is that two lists, or structures, have their corresponding parts noted in the data structure, without regard to the relative sequence of their parts. A zipper-list facility would permit such relationships to be taken note of and kept tract of, despite other developments in the file.

Now, computer people I talked to thought this was a pretty stupid thing to be bothering with, especially since in the glorious realms of list processing you are allowed to tie your data in knots every whichway. Why bother with monogamous one-for-one links when list processing gives you orgies? Yet this is precisely the relation I think is most basic for tasks of creative consideration, although there are many variations. If two things have corresponding parts, whether they are symphonies or machines or individual drafts of some particular book, we need some way to notate these correspondences for computer storage. An item-for-item parallel listing structure is therefore appropriate.

What to call these things is puzzling. Parallel lists? "Parallel" has difficult connotations. "Zipper list" is not bad when the sequence of elements is significant. But when we are talking about correspondences between knotty wholes, such as machines or philosophies, we need a nonsequential term. The term "collateral structures" seems both sufficiently expressive and neutral. I define two data structures as *collateral* if at least one has some sort of private integrity, and the other holds information relevant to it.

The more I thought about collateral data structures, the more interesting they became. First there was the idea of looking up textual parts of one thing by corresponding textual parts of another—a zipper-list pair. Then there was the idea of having alternative versions of the same thing, whose correspondences are kept tract of. More zipper lists. Moreover, the indexing of a list by a list can be extended indefinitely.

Another lead appeared. It might be very useful to have pure text stored in one stream, without codes or irregularities outside the character set, and have its formatting and indexing in other streams. Parallel pointers from outside differentiate the data. More zipper lists. Indeed, it could be useful for many nontextual purposes to store data in one stream and related metadata in another— "a little something on the side."

These different leads seemed to be converging.

Another feeling I had was that there was something basically wrong with the way displays were being programmed and designed. Sutherland's original Sketchpad program seemed to have been set in concrete in the design of various machines, on top of which the conventions of input and output provided many, you will excuse the pun, stumbling blocks. We seem to require *n* buffers, all different; various stacks of niggling errands; and still attention from a main computer. It seemed to me that there must be some way to simplify things.

OPULENCE

There is no time to discuss here the way this has worked itself out, but the

comeuppance has been a rather unusual system design, now ready for implementation. Out of the original desiderata has come an unusual underlying set of structures, conventions, and peculiarly integrated design tricks.

Built around a minicomputer, it should permit the animation of dynamic display without flicker, even as the data rolls or changes; the definition of indefinitely many hypertext types and *connections;* and ways of defining new windows into mind-space. In other words, an opulent input-output machine, sufficient for the richest user activities we may want to program. Its main purpose is to let us create, write, and experience more knowingly, and wander freely through the multidimensional realms of hypertext and hypergrams. But it is hoped that the fundamental design may be extended to provide windows into mind-space of any complexity.

END

If we are to move toward anyone's computopia, or even the simpler goal of a more human and humane world in which computers are prominent, what experts call technically necessary will have to come under close public scrutiny.

Doubletalk and silly press releases have done their damage. The public has been told what experts think the computer should do to them; now the public is down on computers and, by these lights, rightly. It is time for a new accommodation.

And we who have known enough to do so will have to stop fooling the public. To insiders, the computer is not just a tool, but a costume we wear when we want to further our own ways of doing things, much as Bugs Bunny masquerades as a tiptoeing treetrunk. We have gotten away with this long enough. Ordinary people will have to learn enough about computers to evaluate technical assumptions critically, as they already are in the politics of automobiles and garbage.

Computer teaching is an area ripe for public understanding. I suggest that such understanding should begin, not with contemplation of rigidified and sequenced systems, but with an appreciation of the playgrounds and wonderlands for the mind that may now be created. Later we can find what methods of presentation are best, *if any* may be called "best"; but people must understand first the magnificence of computer display and where it can lead.

A sense of awe is essential to work in this area. If there is a failure of awe, you do not understand computer display.

And perhaps that says something about education. For awe and understanding to occur at the same moment is perhaps the pinnacle of the human experience. It is certainly the most important moment of education, if it ever occurs. The two sides of the mind, feeling and insight, are no more separate than the two sides of a coin. Both must be served. Both must act together. How the person feels at the console largely determines what he will learn for good.

I believe one university built in the thirties had a skyscraper called the Cathedral of Learning. That doesn't put it badly. If a cathedral is a place of awe and communion, then our new cathedrals of learning will be our presenting and responding consoles. The architecture of these consoles, and the crafting of their responsiveness and their virtual spaces, is a worthy task.

Starting from general concerns, I have tried briefly to explain one man's obsessions and pursuits, tied together as interconnected ideas rather than merely as an

(77)

enumerated list. I propose two general solutions for a lot of problems: hypertexts for teaching, "information retrieval," and ordinary reading and writing; and, at a very different level, revised programming structures to break the input-output bottleneck for small display computers. But I do not claim that these approaches have unique technical imperative or divinely inspired epistemological virtue. Or that they are the only things worth doing. They are simply among the many, many things that should be tried. To claim otherwise would really be cybercrud.

A CONCEPTUAL FRAMEWORK FOR MAN-MACHINE EVERYTHING

Theodor H. Nelson
University of Illinois at Chicago Circle
Chicago, Illinois

Abstract

This paper proposes a conceptual framework for man-machine everything-- that is, a future condition in which a substantial proportion of our ideas, affairs and apparatus are meshed to our lives by some sort of interactive computer control.

The most general control device, and perhaps indeed the home computer system, will be the text-and-picture display console, and numerous new vivid and pictorial means will be created for presentation and control through it.

Terminology is suggested for a variety of performance features. This is seen as a key subset of the field of fantics, the art and science of presentation to the mind.

MAN-MACHINE EVERYTHING

This paper is not about everything between man and machine, but about man-machine everything, that is, the desirable future condition where most of our information and tasks are attractively and comprehensibly united through nice man-mechanisms. The breadth of possibilities is mind-boggling. But that they are possibilities for the choosing, rather than eventualities to be engineered, does not seem to be clear to people yet. The myth of technological determinism seems to hold captive both the public and the computer priesthood. Indeed, the myth is believed both by people who love, and by people who hate, computers.

This myth, never questioned because never stated, holds that whatever is to come in the computer field is somehow preordained by technical necessity or some form of scientific correctness. This is cybercrud.*

---

* "Putting things over on people using computers." (1)

Computers do what people want them to do, at best. Figuring out what we should want, in full contemplation of the outspread possibilities, is a task that needs us all, laymen no less. There is something right about the public backlash against computers: things don't have to be this way, with our bank balances unavailable

from computers, the immense serial numbers of our drivers' licenses generated by computers, the unstanchable rivers of junk mail sent to us by computers. And it is the duty of the computerman to help demythologize, to help the intelligent layman understand the specifics of systems he must deal with, and to help the public explore the question, what do we want?

Various "professional" approaches to our on-line future have confused us and left us stumbling around. I refer particularly to (a) the field of "computer-assisted instruction," where a computer is often programmed to act like a crabby schoolmarm, coercively leading students around by the nose and chiding them personally; and (b) the field of "information retrieval," where a computer is often programmed to act like a wind-up librarian, sorting new questions into obsolete categories.

But the possibilities are much, much wider, and not to be restricted by the parochialism of "professional" approaches. This paper is not about any branch of computer science, really. It rather approaches the qeustion: "If computers can give us rich services-- and they will-- what do we want? And it supposes that the answer is specific. And that a lot of new terms are needed to cope with the variety of what's coming.

The Home System

Computer fans agree that the home computer is on the way. Soon a minicomputer can be put on a few integrated circuits, and the price will be right-- perhaps a thousand dollars retail before discount pricing. But the question of how we will use it, and thus how it will be marketed, stalls such an enterprise. There are perhaps four models for the Personal or Home Computer. We might want it as a Calculator; Genie; Toy; or Crystal Ball. The question is, what will catch on? We have always had a strange inability to realize what will Catch On Next, though by now we have the vivid hindsight precedents of gramophone, Kodak, telephone, movie, radio, TV, tape recorder, Instamatic, videotape, audio cassette, stereo LP, pocket calculator and so on. But now what?

People are not likely to pay a thousand dollars for a calculator. As to the genie-- something that will open garage doors, manipulate the hi-fi and the model train-- the interface costs are prohibitive. The supertoy idea is also swell, but too expensive for most of us. It is thus inescapable, in my opinion, that the home computer that catches on is going to be a Crystal Ball: that is, its principal function will be as a general-purpose viewplate into realms of digital text and graphics. In that case it may cost more, of course, but in this form the Home System may provide a sufficient and viable basis for a whole new market.

Now, various combinations have been suggested for media of the future, from branching video cassettes to (almost) holograms with dial-up audio. But it seems obvious that when things get sorted out there will be a resolution to fewer things. Just as movie-makers usually do not mix-and-match different forms of output, but stick with sound-on-film 35mm, certain combinations in the grand computer-audio-visual realm will surely predominate. The question is where to cut and combine, what not to bother with, and pre-eminently, what will Catch On. Video is oversold, digital media little known. But I think when the smoke clears our main new medium of the future will be the branching, performing, digital text-and-picture package. Coming over the phone (or other) line to the home system, in pieces summoned by a chain of user choices, it will be almanac, encyclopedia, novel and comic book, playground and travelog and time machine. The Home System will thus be both a Fun System and a Work System.

It is the conjecture here, then, that a Universal Console, a standard text-and-picture demand console, is going to evolve: standard, that is, in its performance and interface specifications, permitting the free interchange of materials. Such a general unit must include graphics refreshment, keyboard input, a selection device, and many service provisions and conventions. There are of course two major ways to do this: as a satellite "terminal" to big computers, or stand-alone.

The usual supposition is that graphics systems need support from a big computer. Indeed, for that case we now have a stunning demonstration that mass computer graphics are practicable, the PLATO system (2). This system, intended for educational purposes but internally quite general, has been developed at the University of Illinois over the past decade by Donald Bitzer and associates. It currently maintains hundreds of highly interactive graphic displays, over a wide geographical area, from a single CDC 6400 computer, and has demonstrated the feasibility of the big computer/terminal approach to mass computer graphics. The PLATO terminal uses a neon panel of 512 x 512 addressable dots with considerable additional vector and character logic to simplify transmissions. While it currently costs about five thousand dollars (that price including other features such as a microfiche projector), it may be expected to be much cheaper later on.

I should stress that PLATO's design philosophy is concerned with other horizons, and its relevance here is as an existence proof, one of considerable appeal and fascination.

The opposite approach is exemplified by XANADU,*

---

* "XANADU" and "Parallel Textface" are claimed as trademarks for computer systems offered by the Nelson Organization, Inc.

---

a system designed over the last decade under private auspices. XANADU is presently under development as a program for a popular minicomputer. However, possible reduction to microprocessors--specialized to the functions of retrieval and display-- is foreseen.

XANADU's basic design criteria were to compress the communication, retrieval and revision of big and fancy files, along with interactive animation, into a small stand-alone machine. The eventual software design, completed in 1972, may best be characterized as a retrieval-and-animation complex handling very large and versatile files. These files, stored on a mix of disk and tape, may be recursively coupled (annotations on annotations on annotations), with couplings surviving revisions; they may have numerous separate data-type breakouts or "enfilades;" and they may be large, being currently defined over an addressability-space of 15×2^{20} elements. XANADU files are subject to extremely rapid revision due to the storage structure, retrieval and edit algorithms; these will, for instance, swap two halves of a large book in a few consecutive disk fetches and writes. Finally, the

system has an unusual display language, DINGO (Display lINGO), permitting interleaved retrieval and animation, and maintaining picture stability while data and picture-parts change. Finally, the design permits incremental roving (see below) in n dimensions of uniform data web, n not related to the number of enfilades.

None of these things seems difficult by itself, or with "enough core" on a big machine; but setting it up to operate on a mini (16K or less) without stopping for breath has been the problem that we believe has been solved in this proprietary design.

XANADU is intended to be the programming substrate for a variety of simple-user front-ends involving complex animation, retrieval and data entry. In addition, because it is designed as a stand-alone system with communications facilities, it is expected to function as a network machine simpliciter, nothing additional being necessary to communicate text, pictures, or interactive animations between XANADU sets. We intend to create a standard XANADU file transmission protocol for all varieties of text, pictures and so on. This involves a rather complex range of mating conventions among files and programs and data, including preambles with faceted data classifications, and default conventions of program and data, so that, e.g., two-dimensional graphic picture-lists can be piped through three-dimensional display programs with

the missing features assigned the proper arbitrary values. Finally, our approach to security involves systems of criss-cross integrity checks to prevent falsification and counterfeiting, a problem that perhaps has not received enough attention for library systems of the future.

One hope is to promulgate this system with sufficient force--e.g., widespread licensing and p.r.-- to create a de facto standard for extremely intricate files. The main problem is, of course, how to enforce standardization in a field where intentional destandardization is a universal lowdown trick.

In any case, since the undemonstrated XANADU system has met much incredulity, we will merely assume here that if this one doesn't work there will be another one of comparable breadth, and talk about what ought to be.

Interpenetrating Screenworlds

We all agree that one way or another, a heyday of computer graphics is coming, and that uninitiated users-- let us call them "simple users," as they may not necessarily be naive-- will use them. But it seems to be supposed by many that the simple user of graphic systems will still have the same psychological environment of today's computer user: he will "call programs" and employ "terminal languages," or at best make selections from uniform-looking columnar menus. In other words, there will still be explicit user-invoked transactions and transitions among data and programs. A little thought may reveal that this is neither desirable nor necessary. We want to be able to roam across boundaries, to call things from one place into the windows of another. Thus tomorrow's sensible graphic systems should permit merged graphic composites-- two-dimensional tapestries or three-dimensional scenes that may be selected and blended from among available graphical and program structures, and roamed over freely by the user.

This suggests that a pre-literate child, for instance, could guide his display screen down a carnival midway with a joystick, turn to watch a cartoon "juggler" do tricks with numbers,

and then, if interested, guide his screen through an entrance into a "circus tent" where the number tricks continue. Or an adult, roving on his screen through explorable views of Stonehenge, may branch from twinkling screen-markers to the many theories about it, and thence to the books and articles expressing these theories--all the while he still explores, and searches out relevant angles in, the three-dimensional model of Stonehenge still on part of the screen.

All this, of course, is not in principle difficult at the technical level: it simply involves certain types of pre-programmed animation among selected parts of complex graphic and text files, and the ongoing selection and retrieval of further files--if these files have been designed for proper graphic meshing and animation.

Such screenworlds can be created for the wholly computer-naive. The more sophisticated user, of course, should be allowed to move with equal freedom through graphic tapestries opening not only into performing graphics and text, but other services and structures as desired.

(89)

The Vending Machine of Ideas

The data conventions of a Universal Console system, then, will allow the interpenetration of contents; for instance, the juxtaposition and inter-framing of graphics, windowing between graphics files, and selection mechanisms among them which include the showing-through of jump markers and other advertising for materials available.

If we call a graphic environment and its rules a "screenworld"--whether a tapestry of drawn data or a set of simulation programs--then this many-ported visual (and calling) access between them creates interpenetrating screenworlds. The advantages of such explorable graphic mosaics should be obvious: roaming over them will be like perusing the Sunday comics (or Ray Bradbury's Illustrated Man), without getting lost, while we remain always in a vividly comprehensible setting. I expect that editorially we will be laying out such tapestries and scenes like magazine spreads.

The question is, what does the human mind want? Given the possibilities of digital exploration, what systems will be best for scholarship, learning, creativity and fun (all closely related)? What are the cleverest and best unifications? We have

yet to find out these answers. But to suppose the desirable systems resemble "instruction" or library searches is hugely premature.

We will probably want a variety of things; these may be grouped loosely under the headings of "responding resources" and "hyper-media." By responding resources I refer to the kinds of things computer people usually think of as "useful programs"-- JOSSes and simulators, timetables and typesetters and so on: services and facilities and programs. Hyper-media*

---

* "Hyper-" here meaning "extended, generalized, and multi-dimensional"-- roughly the mathematical sense.

---

are essentially prearranged presentations without fixed sequence: animated, branching word-and-picture bundles. These include branching and performing graphics, as already mentioned, and branching or performing text, or hypertext.

A few simple examples should indicate the potential power and usefulness of hypertext. Consider first the simple case of quoted material in writing. When we see an interesting quotation in a block of text, it would be nice if we could ask to see it in its original setting, and have the present surroundings fade into the original surroundings of the quotation. We could

read in the original to satisfaction, and only then return to the original setting in which we saw the quotation. This quoteback feature may be thought of as links from quotations to their sources, that we may jump along.

In another application of simple hypertext, many of us long to be able to follow news stories over enough time and detail to transcend the plainly misleading headlines-- but can't, given the existing structure of news media. Hypertext could make it possible. When authors and editors are given the ability to create such discrete jump-links, the character of writing should change dramatically. The potential strength of such new forms of writing can only be surmised at this point, but it should be considerable.

All this was seen by Vannevar Bush in a classic article (3), but what he really said has been largely ignored (4) and the ramifications of this approach--hypergraphics, hypercomics and so on-- have scarcely been touched (5). As with the movies when they were first introduced, most people are having difficulty visualizing the possibilities (6). We may summarize some interesting conjectures on hypertext, or branching text structures (7)[illegible].

read in the original to satisfaction, and only then return to the original setting in which we saw the quotation. This quoteback feature may be thought of as links from quotations to their sources, that we may jump along.

In another application of simple hypertext, many of us long to be able to follow news stories over enough time and detail to transcend the plainly misleading headlines-- but can't, given the existing structure of news media. Hypertext could make it possible. When authors and editors are given the ability to create such discrete jump-links, the character of writing should change dramatically. The potential strength of such new forms of writing can only be surmised at this point, but it should be considerable.

All this was seen by Vannevar Bush in a classic article (3), but what he really said has been largely ignored (4) and the ramifications of this approach--hypergraphics, hypercomics and so on-- have scarcely been touched (5). As with the movies when they were first introduced, most people are having difficulty visualizing the possibilities (6). We may summarize some interesting conjectures on hypertext, or branching text structures (7).

Conjecture 1: we've been speaking hypertext all our lives and never known it. The Tinkertoy structure of thought is inherently parallel, but must be conveyed on the linear conveyor belt of speech. Cross-citing the connections, by intonation, self-interruption, pushes, pops and cross-reference, has always been a day-to-day problem.

Conjecture 2: there has been pressure toward hypertext since the written word began; consider the footnote. Hypertext is immanent in any attempt to put text into man-machine systems, and is certain to emerge no matter where we begin.

Conjecture 3: the interconnective structure of hypertexts will gravitate toward the real structure of the thoughts expressed.

Conjecture 4: understanding of complex relations will come to the hypertext reader through traversal in different directions-- like learning your way around some complicated piece of architecture.

Conjecture 5: hypertext will be easier to write. This is because rather than deciding among expository and transitional structures, the writer may use them all.

It is hard to guess what the forms will be. I suspect that hypertext need not be generally sequential or hierarchical, or general-to-specific in structure, or have obligatory caienas or sequences for the reader to traverse. But these matters remain to be seen.

Prestidigitative Publishing

To make these things possible, then, the Universal Console must be complemented by a range of meshed services: by central feeder machines (large or small), forwarding message and graphic complexes between consoles, and serving up prepared materials. We may call this latter "prestidigitative publishing," involving as it does both the rapid motion of digital data, and the supplying to screens of material that may be controlled "like magic."

The general-purpose text-and-graphic console may thus plug into libraries, explanatory and teaching complexes, literary and entertainment clusters. But we may also expect the basic console to be merged into complex control systems, with a variety of sensors and effectors. (Therefore the data conventions

will have to cover a much wider base, expandable to all possible input and output modalities just as ASCII is expandable to all possible alphabets.) Many modalities may therefore in principle contribute-- text, diagram, video, feelies and smellavision and whatnot; but most basic will be text and pictures, for these will be able to come or go among standard systems.

PSYCHIC ARCHITECTURE

I can now state what I believe to be the central problem of screenworld design, and indeed in the design of man-machine anything-- that is, psychic architecture.*

---

* "Psychic" is here used for the dynamics of feelings and ideas, as distinct from "psychology," whatever it is psychologists study; "mental, as distinguished from physical and physiological." -- Funk & Wagnalls Standard Dictionary, 1960 ed.

By the psychic architecture of a system, I mean the mental conceptions and space structures among which the user moves; their arrangements and their qualities, especially clarity, integration and meshing, power, utility and lack of clutter.*

---

* Similar criteria are also considered in (8).

It should be noted that these notions are much like those by which we judge regular architecture, and indeed the relationship would seem very close. An architectural grand design-- say, of a capitol building-- embraces the fundamental concepts a user will have to know to get around: main places, corridor arrangement (visualization

and symmetries), access structure. These concepts are the very same in a screenworld or other complex man-made virtual structure: main places, corridors or transition rules (and their visualization and symmetries), access structure. It is a virtual space much like a building, (though not confined to three "normally" connected dimensions), and susceptible to the same modes of spatial understanding, kinds of possible movement within, and potential appreciation and criticism.

The orientation problem is in both cases-- real building and screenworld-- immensely important. Because there is no "natural" structure to fantic space (see below), as there is in our 3D world, great care must be given to maintaining the user's clarity of mind. Especially for this reason, the fantic space should have a Grand Design-- an overall shape easy to remember and visualize in some way. I think that there is art to it, that it is not all "human factors" and reinforcement schedules.

I propose the term "fantics" for this newly-structured but very old realm, the art and science of presentation, especially to the mind, sometimes to the hand. I derive the term, like most English "fant-" words, from Greek phainein, show, and its derivatives phantazein, render visible or present to the mind, and phantasia,

appearance or imagination. The related English terms fantasy, fancy, phantom and phantasmagoria (succession of zooming images) also contributed to the forming of this word. Two specific senses of the word phantom should be mentioned. The word "phantom" is used in the graphic arts for diagrams which show opaque things transparently, and also in medicine, for a victim's feeling that a lost part of his body is still attached. (9, pp. 63-70.)

These usages suggest visualizations and physical sensations that come and go. This draws together a cluster of matters I think belong together: showing and presenting things, visualization and kinaesthesis, and alternate ways to structure them in man-information systems.

Thus fantics. Computer graphics will be its principal mechanism, but not its center. Its center is the communication of ideas and thoughts, whether they be facts, poems, or body gestures translated electronically to complex happenings on another planet. Inventing the best presentational media from among the remarkable options now requires our close attention.

Fantic Unification, Constructs and Fields

By "fantic unification," I mean tying things together in a central presentational or control structure which unites them conceptually. Example: several wing-flaps of an airplane are united in its control yoke, a crescent on a rod which may be both turned and moved forward and back. The airplane's flaps do not individually correspond to a desired effect, nor do the combined movements of the control, necessarily; yet this integration provides a convenient unified "feel" to the pilot.

It is in much the same way that we unify things in all presentational modalities-- in writing, in diagrams, in movies or whatever-- creating structures, organizing principles or unifications which have an integrating conceptual character. Often they may have a fictive or not-quite-real component, yet this fiction may contribute some kind of clarity or simplification, allowing the mind more neatly or conveniently to manage information.

A fantic construct, then, we may define as a virtual reference structure used to help imagine or handle ideas and things. It may be added to subject matter or somehow put into a presentational or manipulative system. This concept of fantic construct, then, extends from sequential organizations and headings of text to grossly artificial mnemonics. A fantic structure, however, is the

structure of a presentation or presentational system, whether experienced by the user, intended by the designer or discovered later on by somebody else, or an abstraction never suspected at all.

A fantic field is the fantic structure of a complete and closed presentational, manipulative and/or conceptual field system, within which complicated things may be shown or handled. Thus communication media, such as Sutherland's Sketchpad (10) and the "Parallel Textface" of the XANADU system (4, 8).

books and radio, are fantic fields, but then so are complex interactive screenworlds and work-systems, such as

Fantic controls, then, are any controls whose correspondence to the realm affected is restructured or mediated by fantic structure, whether by fictive fantic constructs, integrated transpositions, or some other form of conceptual combination or rearrangement. Thus wands and puppet-yoke controls, and "virtual gloves" with which we feel inside a display-space, are fantic controls, but brakes and gearshifts are not.

The overall suppositions are these. 1) We are now passing into an era where the structure of objects themselves is less important than it used to be. Not just using or hooking into objects, but structuring the perceptual and conceptual field of the user interestingly and usefully, is the problem. 2) This is the same as the general presentation problem, that of creating presentations in written, audio and other media. Thus we unite

with writing, theater, movies and plastic and graphic arts: organizing for presentation to the mind. The problem is aesthetic as well as cognitive and functional. The aesthetics are important, and, if not inseparable, should not be separated. 3) The principles of psychologically reorganizing receptors and effectors in complex man-machine systems are the same as those of organizing thoughts and other intellectual materials for presentation to the mind. The most basic principles are making things look good, feel right, and come across clearly. Perhaps there are special-case principles, like those offered by learning theory and "human factors," but clarifying their correct range of applicability is essential.

It should be realized that it is not only screenworlds to which these criteria apply, but any media and arbitrarily-structured entities to be presented to people. (For example, if someone were to develop a "hyper-music" with alternatives among which a user could move, it would presumably be subject to these criteria.)

Performance Values and Virtual Spaces

We need a general terminology for the performance features and special effects that we are going to see in the coming years. Unfortunately, because of the variety of devices, modalities, subjects-matter and professional specialties touched, such a common vocabulary emerges only with difficulty. If we concentrate on aspects which are independent of particular areas and subjects, we obtain some generality, if at the price of occasional vagueness. The following terms have come about from rather detailed considerations of possible screen performance techniques, but a wider generality is intended. It is hoped that the following language applies regardless of what we are showing or controlling, or how.

A number of performance features or "special effects" are desirable; we may call these "performance values" (cf. "production values" in films). Many performance values may be turned inside out, and described as if they were places and events.

Fantic space. Space-like structure of acessible text, pictures, animations, etc.*

---

* Cf. "filmic space," the virtual space created by intercutting different shots in the movies.

A space may be one-dimensional (plaintext), two-dimensional (a tapestry) or three-dimensional (a scene or object) or even higher. If it is not regular, but consists of two-(or more)-dimensional zones attached by discrete connections, I suggest the term funny-space (which can be suffixed with the dimensionalities of the zones). Note that the fantic space of the contents joins with the fantic space of the viewing system and controls; the result we should perhaps call a grand fantic space, the overall space the user perceives, thinks about and moves in.

(Note an inversion here: the concept of fantic "space" makes a user's viewplate or other sensorium a moving vehicle, rather than a stationary place to which data are brought. This is as it should be, taking the wider view.)

Fantic contents. The contents of fantic space, as a system of arranged materials and potential performances and events. As in modern cosmology, the space is defined by its contents. Fantic tissue: the connective structure of fantic contents, particularly of their interconnections and transition arrangements. Data web: the data structure which underlies fantic contents, not necessarily homologous or proportional to Them.

Direction: transition gradient in fantic space or tissue; may be presentational or psychological or content-based; may be mapped in visual analogs. Roving: moving through a fantic space. Jump: discrete step between discontiguous places in fantic space. Choice point: place where a choice may be made; e.g., one displaying a menu of jumps. Juncture: joining-place of two continua of fantic space, contents or tissue. If presentable, it may be a choice point involving continuous alternatives, e.g. paragraph-beginnings which continue off the screen.

Rover: a movable place-marker denoting a currently accessible location in fantic space. Jump-set: set of things to which one may jump at a given moment. Fast track: a set of rovers constituting a jump-set; that is, a bunch of markers you can move individually and keep homing to. Jumpstack: a stack holding addresses, in series, of jumps to be undone by a RETURN function. Jumptrack: a currently active system of recorded jumps, all remaining accessible under some scheme, and constituting a fast track.

Border: notable division between spaces, places, media, works, services or facilities, fantic fields, etc. Note that they may connect only at a few points. Opening: local access to another linked place.

Portal: opening permitting full movement of a rover.

Window: opening permitting access but not unrestricted rover movement. Border station: official portal (crossing- or entry-point) on a border. Tunnel: quiet portal between spaces, places, etc.; border is not seen. Customs: markers, crossing protocol and/or restrictions at a border. Crossing zone: area of parallel or other multiple access across a border. Seamless tssue or web: tissue or web having no break or sharp discontinuity in performance at a border, or simply no borders.

Link: connection between two points, as between rovers and placemarkers, which may serve as a jump from one to the other; connector between ends of an opening. Coupler: facility permitting a link to be made, or data structure resulting. Prehensile coupler: one coupling into a file that has not been forewarned. Multicoupler: facility permitting multiple links to be set between parts of two entities, thus establishing a sort of correspondence between these parts; the resulting data structure (also called "zipper list," 11). Collateral structures: entities linked by a multicoupler.

Mooting system: system permitting complex alternatives to be studied indecisively, e.g. by use of collateral structures. Creativity system: mooting system with design (or text, etc.) facilities for creating complex entities indecisively.

Visual Orientation Devices

Setting aside some of this intended generality, let us consider the visual modality, and screen-tricks to keep the user oriented during complex screen transitions.

Orienters: pictographs that show a user where he can go spatially. Rose or compass-rose: pictograph showing possible "directions" in the space. (Map: diagram showing the structure of the space, and perhaps the user's current position.)

Ticklers: pictographs or messages that tell a user what he may do next. (Not sharply distinct from orienters.) A maplet, for instance, is a fractional map that tells immediate alternative moves, and possibly which way to some sort of "home;" a blurb is a writeup or title of something that may be gone to; a jump-marker is an element you point to to jump. Function box or function rose: pictograph indicating alternative functions which may be chosen. A menu, of course, is a textual listing of alternatives.

Note: for fast-roving performance on slow-filling screens, it is desirable to put the ticklers up first in each new screenful, permitting the user to jump or move at once.

By the careful crafting of media and screenworlds, using these devices with as much consistency and attractiveness as possible, we may encourage well-oriented mobility in our fantic systems. This is the point.

A FANTIC AGE

A variety of media await us, and which ones make the "best cuts" remain to be seen. (Conjecture: they will resolve to a very few.) The major mixable options I wish to point out here are: "hyper-" (nosequential); text; pictures (2D or 3D); animation; overlays; complex roving; complex coupling of structures; interpenetration of presentational tissues; linkage of video and movies; linkage of special control yokes; linkage of special forms of viewing (3D, smellavision, etc.).

Except for the more expesive viewing situations-- film, smellies and the like-- all these may be presented in the Universal Console, and those which cannot may still be comprised within a generalized data and transmission structure built around such a Universal Console.

Human Intercourse in a Fantic Environment

The spread of computer graphics and the prospect of a Universal Console will mean that computer systems may furnish backgrounds to human intercourse as varied as the blackboard and drive-in movie.

Till now, graphic systems have been built with an organization-based frame of mind. They have involved sitting

in office-type chairs at desk-like furniture having screens and keyboard perched or attached. Home models, however, may well be built into lounging chairs, conversation pits, children's furniture, or even into bedlike environments. I am particularly interested in a backpack design for portable wear, one which would mirror the CRT image into a concave transparent faceplate or visor. (This is a simplification, of course, of Sutherland's exquisite helmet. (11))

The active use of such systems in conversation should make possible whole new depths of communication, both factual and ideological. Engelbart's recent work has shown the power of text CRT systems to enable people to work together over the same materials. Perhaps this ability can be extended to miniature portable systems delivering graphics as well, e.g. where people can exchange graphics or text through a quick umbilical connection between backpacks. Suppose we can carry our favorite animated diagrams and reference works around, and exchange them and talk about them and manipulate them together. Does it matter how turtle-like such portable Houses of Intellect might seem? May we not actually come to understand each other better?

Should Systems Talk?

It seems to be a general view that computers should talk back. From Weizenbaum's ELIZA program to HAL-9000 of Kubrick and Clarke's 2001, there is a constant sense that "of course" we want talking computers. And while the problem of how to get them to talk back is investigated every which way, the question of whether computers should talk back seems never to have been examined, let alone posed. Given that there are far simpler methods of commanding system activity-- e.g., the light-pen thrust-- conversation is by no means necessary. Moreover, it is not obvious that people will enjoy mechanized conversation; rather, it may be offensive, alarming and a tribulation. Mischievous programmers get a kick out of writing programs that pretend to speak and understand, wise-guy programs that identify themselves as "I, the computer," insist on being talked to by typed input, and are full of snappy replies but don't really do much with the input. The capacity of such systems to offend and annoy has not been sufficiently recognized.

Now, obviously there has to be a method of sending new information to the user, and so computer sentence generation is unavoidable. But that does not mean a system has to be a smart aleck, or, indeed, to disguise the exact method of its sentence

generation procedures.

What I am getting at is this: we should have standard ways to introduce systems, so you know what kind of an entity you're communicating with. Who wants to be the goat of the Turing-test? The user is entitled to know if he is typing into a real sentence-parser or just a keyword trickster program, or what kind of a mix.

It is infuriating to have a program pretend it can understand you and then fail to parse eight consecutive input sentences. There ought to be a law against wasting people's time with this sort of silly program.

More generally, the spread of consumerism means people want to know who they're dealing with. As we program on-line systems that will involve innocent people, we had better think hard about ways to get the system's cards on the table (or heart on its sleeve)-- and play neither coy nor god.

Victorian Remarks

I think that the Grand Corpus of our written heritage, chaotic and individualized as it is, is a precious substratum of our world. The new age of hypertext and hypergraphics should build on this tradition, rather than mush us into committee authorship and indifference to the past. In forging toward the Screen Future, and the creation of screenworlds we will love to live in, let us remember and esteem the traditions, scholarly mechanisms and arts that have worked so far, and build on them. And we must begin to worry about the problems of privacy, access for everybody, "what gets kept?" and dangers to the corpus once it is on-line.

Conclusion

The principal development foreseen here has been the extension and unification of the writen word-- our libraries and leisure reading-- and various forms of pictorial information and control through generalized consoles and interfacing conventions. This will be the apotheosis of computer graphics; but the sooner wer resolve "computer graphics" to a set of techniques at the service of explicit presentational goals, the better.

This paper has presented an ambitious terminology. The objective has been to nail down crucial aspects and distinctions of an entirely new realm of human endeavor and experience, not cluttered with mumbo-jumbo from other creeds. The psychic

engineering of fantic fields--adults' hyperspaces of word and picture, child's gardens of verses--is our new frontier. We must look not to Asimovian robotics and the automated schoolmarm and librarian, but to the penny arcade and the bicycle, the clever diagram and the movie effect, to furnish this new realm.

References

1. Theodor H. Nelson, "Computopia and Cybercrud." In Levien (ed.), Computers in Instruction, the Rand Corporation, 1971.

2. D. Alpert and D.L. Bitzer, "Advances in Computer-based Education," Science, 167, 1582-1590 (March, 1970).

3. Vannevar Bush, "As We May Think." The Atlantic Monthly, 176.1 (July, 1945), 101-108.

4. Theodor H. Nelson, "As We Will Think." In proceedings of Online 72 Conference, Uxbridge, England, 1972.

5. -----, "No More Teachers' Dirty Looks," Computer Decisions, September 1970.

6. -----, "Getting It Out of Our System." In Schechter (ed.), Critique of Information Retrieval, Thompson Books, 1967.

7. -----, "Hypertext notes," unpublished, 1966-7.

8. -----, "Simplicity Versus Power in User Systems: The Text Case," unpublished, 1972.

9. Paul Schilder, The Image and Appearance of the Human Body. Science Editions, paper, 1964.

10. Sutherland, Ivan, Sketchpad: A Man-Machine Graphical Communication System. Lincoln Laboratory Technical Report 296, 1963.

11. Theodor H. Nelson, "A File Structure for the Complex, the Changing and the Indeterminate." In ACM Proceedings of the 20th National Conference, 1965.

Data Realms and Magic Windows

Theodor H. Nelson

This talk will contain approximately one thought, which I think is the optimum number for an after-lunch address. None would be too few, though it is a number often chosen, and two thoughts would get in each other's way and become confused.

I want to talk about computer programming for business, a subject in which some of you are involved, even interested. But I will talk about it as an outsider, perhaps with some of the inspired vagueness which is permitted on these occasions.

First, however, let me explain the point of view I am starting from. I have for some time been engaged in a project which is intended, to put it modestly, to be the library and publishing system of the future.

Now, I am going to have to talk about that for a little while—what that system is about, and what it is attempting to do. I thought I would talk first about the assumptions and then about the actual design, but these easy distinctions do not work out in practice. The early assumptions are embedded in a philosophy, but that philosophy has gotten so deeply intermingled with the developing final design that there is no good starting place. Here is most of it.

We will have to stop reading from paper. The trees are running out, the distribution and disposal of documents cost too much, and we can no longer find what we need in an increasingly fragmented world of information.

Though most manufacturers steadfastly refuse to understand it, the next big computer market is at the consumer level. (Thousands of Altair computers, priced under a thousand dollars each, have been sold this year already.) But consumers will rarely want to program. The next grand market in the computer field will be variety of consumer turnkey products—first games, then simple interactive bookkeeping and retrieval programs, such as for phone numbers, checking accounts and magazine subscriptions.

Then will come the libraries. I think it is obvious that library service is the frontier of consumer applications, and that few computer people see what this entails in all its magnitude.

A few years from now many people—later "everybody"—will have a computer terminal. But not just a slow old printer: a fast, high-performance vectoring CRT, which is cheaper to make and simpler to use. Writings can then be stored in full-text digital form, and readers will have highly responsive consoles—preferably of the vector-display type, permitting the text to move smoothly at any speed under throttle control.

But merely supplying text to the reader's screen is hardly enough. Not only must your library-CRT system supply vast quantities of text at any usable rate of speed; it should also bring capabilities which are not now available in the best libraries. More about that later.

User actions must have immediate results, preferably under three seconds. This means computer capability in the terminal, as well as some local mass memory. Now, computer mainframes—the processor chips—will cost very little; likewise core (or RAM chips) will be cheap as dirt; but there will always be fast expensive memory and slow cheap memory. Thus for cost reasons the common notion of "core being cheap enough so we can have everything in all the time" is forever absurd. The system must come to grips with this great truth, in terms of managing large multilevel stores on small machines.

Similarly, "whole libraries" cannot be stored at the user's terminal, but must be rapidly accessible through communication lines on demand. It follows that the system will use computers in a net. These need not be big computers, as they are to be used only for storage and distribution—communications processors. Thus we come to the notion of a network of minicomputers. But the simplest and most elegant approach is to have all units on the network—user nodes and communications processors—use essentially the same program.

This is the fundamental idea of the system.

Now, when you say "library," computer people say, "Ah, yes! The Encyclopedia Britannica and Moby Dick." But we need to do more than rapidly deliver linear text. It is my working belief that linear format is not the best organization for anything, but is forced on us by the constraint of printing on paper. The real use of information by people, for a spectrum of uses involving everything from CAI to IR, will be best served by new non-linear forms of writing, generically called hypertext. I have discussed this elsewhere. *

* Theodor H. Nelson, Computer Lib (published by the author; available from Hugo's Book Service, Box 2622, Chicago IL 60690). See esp. pp. 65-70.

This is not clear to many people, because they are so used to linear writing—but our use of writing is rarely beginning-to-end linear, and neither are the thoughts that have been pressed into linear sequence.

Indeed, ordinary writing (ignoring illustrations for the moment) is only quasi-linear. While the author gives us a principally linear string, there are also little jumpouts—usually footnotes, sometimes "boxes" and headlines—that break out of linear format and permit the reader to pursue matters he guesses to be attached there. The succession of points and cross-references in a document are seldom linear. Moreover, we do not ordinarily read in linear sequence; and after we have read something, what remains in our minds is rarely linear either.

But to recognize this is just the beginning. When we acquire a high power non-linear text capability, considerable research and inspiration will be needed to find the best forms of non-linear organization for writing.

Very well; we must supply fast-access, non-linear library services. But even that is not enough; we want also to provide users with various helpful features for their own reading and note-taking. For instance, placemarkers (to work like bookmarks), allowing instant jump to things a user has previously flagged; and software allowing users to make marginal notes.

This leads to an interesting compound capability. Not only can an author create materials in complex non-linear document form, but a reader can incorporate such a non-linear document in a complex of his own comments, annotations and additions. Now this is a curiously involved outlook. Such compound linkages are not like anything seen usually in the computer world. But that doesn't mean it has to be disorderly. While the notion of infinite cross-coupling unnerves some people, it may be kept orderly. Whatever the complexity of the author's original version, and whatever the complexity of the user's linkages and even modifications, an orderly systems approach can keep it all systematic and sensible.

The fundamental way of keeping all this straight is with a concept implemented basically in the system. Namely, every document has an author, who owns it and is the only one who may change it. Each reader's version, containing links and modifications of his own creation, belongs to him—except for the parts showing through from some other original. In other words, the concept of authorship is to be scrupulously preserved, as well as the author's canonical version or versions. These provide the anchor for a number of interesting services.

I expect that the roles of writer and reader will remain as they always have been, and so will their motivations: for the authors, money and respect (sometimes fame); for readers, orientation and background information, the answering of questions, the enhancement of general knowledge, self-improvement, the pursuit of interests, entertainment and fiction.

To further this general idea, an unusual system of proprietary techniques has been developed for the storage, linking, retrieval and editing of large non-linear documentary structures. The intention is to use this retrieval subsystem as a basis for a product line, beginning with a simple text editor (slated for 1976), and working up to the full system. We have discovered certain structures and techniques we consider fundamental and highly generalizable. Although our data structures and algorithms were designed for the library network application I have described, it may be possible for this system to have applications in more "practical" environments. For instance, it is possible that a whole package may be marketed as a virtual-memory data-base add-on for regular general-purpose setups.

I have been frequently criticized for saying what this system will do without either showing it or telling how; in this situation I am somewhat like Frank Marchuk, the man who says he has a "laser computer." But I know that a few of you are interested in keeping track of the system prior to its actual appearance, and these remarks are for your benefit.

The system under implementation has several properties of note.

A single system is complete: from the minicomputer the user may access and revise files of any size, provided they are on local storage.

A plurality of units automatically becomes a network when hooked together, with each processor looking to the other processors merely like a terminal that wants access to a file.

"Linked text" may be thought of as the canonical storage type of this system; but properly understood, linked text can be anything. The system holds strings of any length (up to total hardware limit — with no access penalty for length); and tables and arrays, either thick or spares. (An element of an array can be a superlong string.) Between elements there may also be links; such links may be modal, and defined and flagged arbitrarily.

A file, or organized collection of data, is stored as what we may call a hyperblock—a virtual structure of potentially great size whose parts may nevertheless be rapidly accessed. (For optimization in nonliterary applications, there is a certain leeway in the declaration of preferred access methods and classess of path. But this will be discussed at a future time.)

The problem of "data integrity," especially under the multiple ingressions of inquiry and update, has been cited as one reason for needing big computers. This is otherwise settled in the present system. Our approach is that every file has a current hardware captain, which can be any computer in the network. Only one computer is at any time the captain of a file, and it will process no more ingressions until all vital updates have been made.

Certain postponements of update are possible. Catching up with these is background to the foreground features of display, retrieval and edit. The rapidity of certain astonishing events is actually sleight-of-hand. Certain "instantaneous" complex events are actually not complete at the moment they appear to be, but they are in an update pipeline, and their far-flung ramifications are actually consummated in a system of ripples propagated through the file as a background task. The system is thus slightly marred by, ahem, catchup stains.

In describing this system, I have left for last the parts I find most interesting, those which lead to the one idea I wish to present. Let me describe several unusual facilities the system is intended to have for non-linear writing.

One is the facility for collateration, or the multipoint linkage between two text units. This resembles (and can be used like) marginal notes or commentary between two texts or for various other purposes.

Another is the facility for quote-windowing. This is a very interesting and unusual capability, around which I expect some remarkable uses to evolve.

The idea is that one piece of text can have a quotation from any other piece of text in the system. But it is not merely a quotation. The quoted text is actually still sitting in its own file; the user reading the quote is actually at that location in the quoted file, as well as still being at that location in the quoting file. In other words, a directional splice-point exists from the quoting file into the quoted.

The main purpose of this is, of course, so that the reader may, should it please him, transfer into the quoted file and stay there.

Scholars should find this useful for checking sources. A researcher looking for something may use it to go from a vaguely relevant source to a more exactly relevant one. The quote-windowing facility also provides a way that an anthologist, writer or correspondent can provide his reader with instantaneous entry to a number of things he thinks the reader ought to see. And, finally, it suggests an interesting solution to the copyright problem: since the quoted material is never put into another file, it can be a simple matter to count and pass on royalties everytime that file is entered by a user.

While this windowing function is most easily described for text, it is intended to apply to graphics as well. Thus one "picture" from the library may contain a "window" into another picture from the library, ad infinitum. These excerpted pictures may even be animated.

Very well. I have described an unusual system now undergoing implementation, both to show the point of view I'm starting from, and to initialize certain ways of thinking. The application of this Windowing philosophy to business programming may turn out to be rather interesting.

Now recently, on the basis of this background, I had cause to think about The Business Problem. The stimulus was the June 30, 1975 issue of Business Week, which offered an "executive briefing," whatever that is, on "the office of the future," whatever that is. The Office of the Future as they described it—especially as a marketing concept advanced by a major manufacturer for mechanized typing pools—was so appalling, bleak and absurb that no thinking man could contemplate it without a shudder. *

* Business Week, "Putting the Office In Place," June 30, 1975. pp. 56-70.

The general view was this: the Office of the Future will be sharply divided between peons and executives. The peons will be typing clerks who have been taught to work such

"word processing" gizmos IBM's mag tape and mag card typing machines. The peons will be herded into central typing pools, where they will finger in what the executives dictate. Only a few lucky secretaries, the real smarties, will escape this fate. Lucky them, they will get to file the paper that the Executives dictate and the Peons finger.

Now, this is so wrong and awful it needs no direct reply. Let me speak instead about how things should be. Let's talk about how computers should work, and call it ten years from now to keep your minds off implementation problems.

Look, isn't it obvious? There should be a general system of some kind, and everybody should have a display terminal. Your job determines what you may or may not see or change.

For each job there are clear and simple interactive programs—let us call them "workfaces"—for all the things to be done. One person may have access to dozens of workfaces, and the choice, perhaps, of which workface he prefers for a given problem. The workfaces are clear and simple. Some are more general than others, like the workface that helps you type in a letter or a memo.

There is no paper anywhere, except for letters which arrive from elsewhere or those which finally leave the office. All putters-in of information have clear and simple interactive programs to help them. Moreover, all info-dippers, too, have simple and clear interactive programs to help them. But this distinction, long with us, will erode: probably every job definition in the future will include input, viewing and revision.

Good interactive programs for everybody will enable and encourage us to use our minds better. It is my belief that people will turn out smarter, happier and more productive in a system where they understand what they are doing, and are encouraged to use their minds.

So much for the preamble. I will now endeavor to express the one thought with which this talk is concerned.

Let us suppose that I am generally correct about what the office of the future should be like. Let us further suppose that a data system such as the one I described earlier—holding vast quantities of stuff on a network of minis with big secondary storage—is feasible. How then can we put this righteous Office of the Future up on it?

Earlier I made a couple of major points about such a future.

First. Central to this view was the notion of future life being organized around big files of text and graphics, stored in distributed form on a network of minicomputers. According to this idea, information may be thought of as making up big "documents" of one sort or another, especially text but also graphics.

Interactive business systems, too, can be defined in terms of "documents." If everyone in the business is to have a screen terminal, written clarifications will be needed for most types of data. Let us consider, then, a hypothetical business system that consists of a long written report on the state of the business—with windows that can show the changing summary figures, status information on outgoing and incoming orders, sales reports and prospects; and so on, for every facet of the business. (We will overlook here the problems of security and restricted access, but these are naturally a part of the problem to cope with eventually.) Here, too, then, we have a virtual space which is also a virtual document, that is, a series of presentations with windows into the data.

Now that document, some large part of which can be text—a sort of standing, up-to-the-minute corporate report has "distances" in it. Relative "distance" in the data structure is the key concept. For instance, the two elements at the ends of a long text string are far apart, and will not ordinarily be held in core at the same time. A customer's name and i.d. number, however, are as close as two things can get; his address and phone number, are presumably at one removed (requiring the following of a pointer), and the salesman's comments on his purchasing prospects somewhat further away (requiring both pointer and lockwords).

Now, it seems to me that things are still too complicated in this field. Data base programming should be no big deal and should be virtually taken care of by the operating system; defining the data types and their access chainings should be about equivalent in complexity to a Sysgen. Now, as I envision it, the data base system would have certain given characterisitics any programmer can learn in a hurry.

Because core is small, the programmer declares what forms of information will need to be considered together. He declares access paths between types of information, such as what kind of thing points to what and what has to be found independently (such as alphabetically or by number). Things are either in core together or not, directly accessible or farther away, and so on.

Second. You will recall that I also said the key to such a data network was the use of staged update. While the system always has instantaneous integrity, it is not always caught up. That is, various temporary mechanism take note of whatever changes the user invokes on the data, whereupon the system replies, "It is done—" but the actual meaning is, it is as good as done. Finalization of the changes in canonical stored form occurs gradually thereafter. But meanwhile, all inquiries see only the current virtual structure, as ingeniously declared in various exceptions and memoranda in the system.

I submit that the way to make this distributed-mini office of the future work is a simple development of the same idea. The data base is to be a staged system. Rather than consummate all the consequences of each input keystroke—an obviously absurd extreme alternative to the batch approach—inputs are accumulated in partial forms. These partial-products are assimilated to the data base and consummated either in the course of events or as demanded.

But demand is in the form of user requests to look at something. As in the previous system, wherever he looks, it appears to be finished.

As the user gets near something, the appropriate crunches begin. The procedures of final data consummation are done in catchup mode, clued by the user's present location. This is rather like burglars peering at a night watchman on his rounds.

In other words, the programmer, in the course of preparing an interactive system or "workface," ascertains what partially-processed data must be further refined for a given user activity; these catchup actions are then cued as the user approaches places where these results will be needed.

Now, this is rather interesting from a computer-science point of view. Places invoke procedures. A user's arrival at specific places in the user space means actions on the data.

Obviously it would be senseless for each individual routine to test where the user is; every routine would have to be called in all the time to see if it were needed. Obviously, instead the general monitor tests where the user is, and invokes the routines called for by the user's postion. But that means there has to be a table of what routines occur at what places.

This means in turn that a table of program cues must be compiled from the total collection of all catchup procedures, so that those affected are appropriately notified by the user's position. This calls for a sort of backwards compiler—let's call it a Relipmoc—that assembles such a list and associates or inserts it in the appropriate places.

Now obviously there is a problem. It will readily be seen that certain contentions will occur. Certain "places" in the

(17)

system, as it evolves, may come to trigger too many separate actions.

There are two major lines of solution. Programming staff may have some political procedure of arbitration that decides among the candidate procedures, such as the arbitrary decision of the lead programmer.

Or steps may be taken to control the user's movement, slowing him down when there are too many things to be done. If the catchup time for a given function is longer than it will take the user to get there, various delaying techniques may be employed which keep the user occupied until he's ready. Indeed, usefully occupied. If need be, we can delay the user with other things—entertainments or reminders of useful sorts—till what he wants is ready. ("Vamp till ready," they called it in vaudeville.) Or we may modify the "documentary" setting, with new writings—preambles, updates, and reminders that slow him down.

(This is an example of one of the most interesting options of the man-machine system: the option of changing the user's experience when the program can't be done as originally intended.)

Third. I have spoken of saving various keystrokes in preprocessed and summary form, as "partial products." In the scheme envisioned here, the keystroke capture programs will update the partial products, and the display will invoke their final conversion and assimilation.

Thus the major decisions to be made at the systems level, curiously enough, are those concerned with the form and content of such partial products. They will of course include sums, transactions (such as the latest news about whether a specific invoice gibed with an incoming delivery), and updates to individual records. I will admit that the weakest part of the scheme lies here: in the practicality of the half-deferred update for general cases. I simply pass the idea on to you, who understand more about business programs.

So much for my one thought and its possible ramifications for business programming. Through these proposed mechanisms we get our keenly responsive, people-oriented office of the future. I assume here, of course, the availability of a high-performance retrieval system roughly like that already discussed, whose principal foreground task is interactive display. I also assume that people's time is more valuable than computer time, and the computer must do its "work" at times most convenient for its users.

Obviously the programmer's task has changed. Rather than living the batch life, or even the usual database life, in a way he now scuttles in shadows. His function is to implement the desired user experiences, not to dictate.

But there is still work for programmers. This consists of declaring the database (roughly a one-shot deal); and, to service the users, the creation of the workfaces. This in turn requires programming to and from the partial products by whatever transformation techniques seem most appropriate.

In all this the user's experience is central. When the techniques must be traded off against user experience, it is the experience that counts, not programming convenience.

It is this curious tradeoff that has made me coin the term "fanties" for the art and technology of showing things to people. I believe this is the correct generic term under which interactive programming falls. I have been criticized for making the term too broad. Yet this is all an indivisible whole, and if we narrow the discussion to a smaller part, such as computer graphics, we overlook both the overall effects of what we do and the options we might have tried. "Fanties" is a single tradeoff domain. We can present the same things on paper, on screen or by ear (or using the techniques of Madison Avenue, as a joke, in a song or a playlet); we can have the user point with a light pen, type on a keyboard, tap his feet or whistle.

These options are none of them right or wrong. It is how we tie them together that is good or bad, clear or confusing, dreary or fun.

Programmers are very goal-directed, sometimes too narrowly. There are different kinds of goal-directedness. A truck driver, or tank corps commander, is given an objective and perhaps a few limitations on what he may do to get there. I submit that where users' convenience is concerned, that is not the right kind of goal-directedness.

Consider the goal-directedness of a stage director: his objective is not merely to move the actors in and out and have them speak the correct lines; that is just the beginning. His objective is to create a gratifying overall experience for the people out front, merging the talents and facilities that he has available. This kind of goal has not single event, with Boolean success or failure as its culmination, like driving a truck to Detroit or computing pi to a million places. The goal is a series of events, each of which can be carried off well or poorly, each of which can gratify or annoy, and which finally result in somebody else feeling bad or good.

That feeling may be the only purpose, as in game programs and computer-controlled artistic events. More often the user's feeling is a byproduct, as in a hospital sign-in system, but that makes his feeling good no less of a goal. Few organizations today are unconcerned with the morale of their members and their clients. If you insult a user, confuse him, make him feel inferior or bore him, good will has been lost that will undercut all the other goals of the system.

The programmer may look at his job narrowly, and pretend he is only driving a truck to Detroit. He may even convince himself (and his supervisor) that making a more suitable nest of code run "correctly" is all there is to be done. But it isn't so.

The interactive programmer should understand that the mere computer is not the real focus of his work. He is working in idea-space. The feeling and the idea of the system to the user emerge from all the parts acting together. The user's idea of what he is doing, and what he is doing it to; the constructs and the spaces he works with and in; the clarity of it all; these are generated in non-simple ways by the program. The computer is a paintbrush for a mural of experiences you are trying to give the user.

Through the Magic Windows we give him, we want the user to see his Data Realms in bright sunshine.

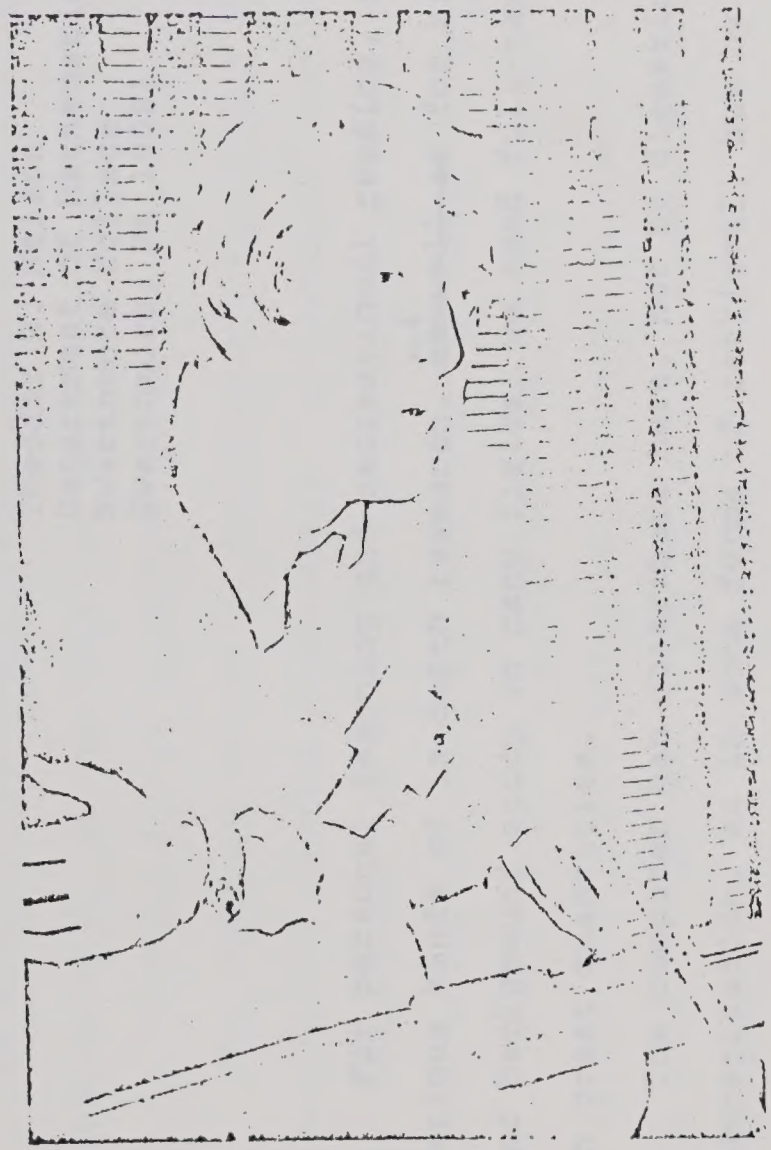

Ted Nelson

. Theodor Nelson describes himself as a "writer, showman, and generalist." Ted is both author and publisher of COMPUTER LIB / DREAM MACHINES.

. . Ted holds a B.A. in Philosophy from Swarthmore and an M.A. in Sociology from Harvard, where he studied under Henry Kissinger, Timothy Leary, and Talcott Parsons. He has worked for Dr. Lilly's Dolphin Lab, taught sociology at Vassar, worked as a professional folk singer, and as a consultant to the Anti-Ballistic Missile System at Bell Labs.

. . He spent several years as proprietor of The Nelson Organization, Inc., which he founded at the express request of the Central Intelligence Agency. Ted is currently a lecturer at the University of Illinois at Chicago Circle, in various departments, including Art.

From ACPA-V Proceedings (Proceedings of the 5th National Conference of the Association of Computer Programmers and Analysts, 1975).

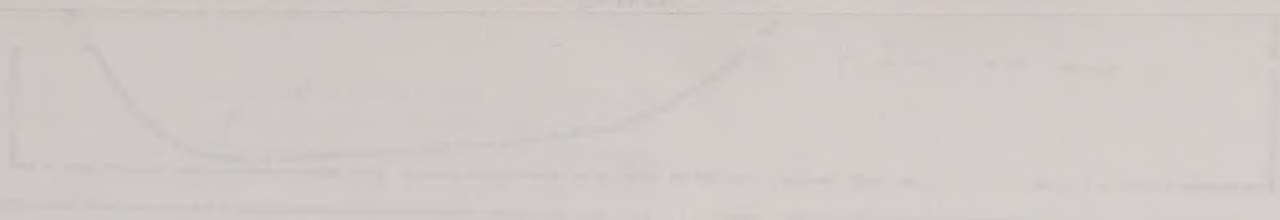

(119)

[for National Computer Conference, 1977]

12 Jan 77

DESIGN OF A TRANSCENDENTAL LITERARY NETWORK

Theodor H. Nelson
Department of Mathematics
Swarthmore College
Swarthmore, PA 19081

For personal learning and recreational reading, for various kinds of in-depth research, ~~as well as~~ and for historical and background study in many fields, we need full-text access in great quantities.

The computer can contribute here, not by digestion and summarization, as in some forms of retrieval, but by extremely quick full-text delivery, and especially by certain sorts of cross-access.

The network proposed here will hold an extraordinarily large corpus of documents for rapid viewing, and for extremely fast cross-fetching along links between them. These links, put in freely by authors and annotators, may be freely used by the reader as paths leading to other documents.

The types of links proposed here carry out automatically certain lookups which have long existed implicitly in ordinary literature: the jump-link (corresponding to the conventional footnote and citation); collateration (corresponding to side-by-side printing and marginal notes); and quote-windowing, whereby the original context of a quotation may be summoned. While other linkage modes are possible-- notably annotated

and "colored" links-- this set has the advantages of simplicity, recursive extensibility and psychological clarity over many removes.

As these linkage facilities extend the notion of literature considerably, we use here the term "transcendental literature" for a corpus of this type,[25] available in this delivery system.

The characteristics of document and authorship are extended to promote orderliness. A "document" may consist of text, links, or both; it has an author or owner. Thus there is no ambiguity of ownership, or danger of the sullying of a work's integrity by someone who is not its author. However, any author may freely annotate or create links to the work of another; these contributions becoming his own "document."

Rapid text-editing facilities are an important feature of the system, as are the storage (if desired) of successive drafts and modifications, whose canonical archival form permits point-by-point comparison among successor versions. Simple extensions will permit string-search and concurrent thesaurus maintenance.

Extensions to graphics and animation are contemplated.

Tele-conferencing capability is of course implicit , enhanced by the ability to link into the greater stored corpus.

Because this is a special-purpose multi-user delivery system, employing minicomputers in a network, low prices are projected: two to three dollars an hour including author

royalties and all basic services, exclusive of communication costs. However, user-written programs cannot be accommodated.

As vital as the technical aspects of the system are its planned accommodations to a free-enterprise economy: notably, that the condition of fast accessibility must be paid for by someone-- called a "publisher"-- and network nodes must be financed through the free-enterprise mechanism, possibly by franchise. Finally, copyrights and royalties are handled automatically within the system.

It is often carelessly asserted that no such rapid full-text delivery system will be useful until it contains "the entire Library of Congress." Quite the contrary: client groups with special interests may justify their own usage for a few million characters of complicatedly linked material.

Most important are the issues of privacy and freedom of speech and the press. Safeguarding these in the ever-growing corpus of our civilization has increasingly technical aspects, but is not a technical problem.